arXiv:1610.03113v1 [stat.ML] 10 Oct 2016

Truncated Variational Expectation Maximization

Jörg Lücke
Machine Learning
Department of Medical Physics and Acoustics and
Cluster of Excellence Hearing4all
Carl von Ossietzky Universität Oldenburg
26111 Oldenburg
Germany

October 12, 2016

Abstract

We derive a novel variational expectation maximization approach based on truncated variational distributions. The truncated distributions are proportional to exact posteriors in a subset of a discrete state space and equal zero otherwise. In contrast to factored variational approximations or Gaussian approximations, truncated approximations neither assume posterior independence nor mono-modal posteriors. The novel variational approach is closely related to Expectation Truncation (Lücke and Eggert, 2010) – a preselection based EM approximation. It shares with Expectation Truncation the central idea of truncated distributions and the application domain of discrete hidden variables. In contrast to Expectation Truncation we here show how truncated distributions can be included into the theoretical framework of variational EM approximations. A fully variational treatment of truncated distributions then allows for derivations of novel general and mathematically grounded results, which in turn can be used to formulate novel efficient algorithms for parameter optimization of probabilistic data models. Apart from showing that truncated distributions are fully consistent with the variational free-energy framework, we find the free-energy that corresponds to truncated distributions to be given by compact and efficiently computable expressions, while update equations for model parameters (M-steps) remain in their standard form. Furthermore, expectation values w.r.t. truncated distributions are given in a generic form. Based on these observations, we show how an efficient and easily applicable meta-algorithm can be formulated that guarantees a monotonic increase of the free-energy. More generally, the obtained variational framework developed here links variational E-steps to discrete optimization, and it provides a theoretical basis to tightly couple sampling and variational approaches.

1 Introduction

The application of expectation maximization (EM; Dempster et al., 1977) is a standard approach to optimize the parameters of propabilistic data models. The EM meta-algorithm hereby seeks parameters that optimize the data likelihood given the data model and given a set of data points.

Data models are typically defined based on directed acyclic graphs, which describe the data generation process using probabilistic descriptions of sets of hidden and observed variables and their interactions. EM approaches for most non-trivial such generative data models are intractable, however, and tractable approximations to EM are, therefore, very wide-spread. EM approximations range from sampling-based approximations of expectation values and related non-parameteric approaches (e.g. Ghahramani and Jordan, 1995), over maximum a-posterior and Laplace approximations (e.g. Olshausen and Field, 1996; Bishop, 2006; Lee et al., 2007; Mairal et al., 2010) to variational EM approaches (Saul et al., 1996; Neal and Hinton, 1998; Jordan et al., 1999, and many more).

Instead of aiming at a direct maximization of the data likelihood, variational EM instead seeks to maximize a lower-bound of the data likelihood: the variational free-energy. Variational free-energies can be formulated such that their optimization becomes tractable by avoiding the summation or integration over intractably large hidden state spaces. Since the framework of variational free-energy approximations has first been explicitly introduced in Machine Learning (e.g. Saul et al., 1996; Neal and Hinton, 1998), variational approximations for the expectation maximization (EM) meta-algorithm have been widely applied and were generalized in many different ways. Variational EM is now routinely applied to basic models, to time-series models as well as to complex graphical models including models for deep unsupervised learning (see, e.g Jaakkola, 2000; Bishop, 2006; Murphy, 2012, for overviews).

For the purposes of this paper, we will introduce our novel variational EM approach using the basic setup of the well-known paper by Neal and Hinton (1998), i.e., we assume one set of observed and one set of hidden variables. Such a setting is best suited to introduce the basic ideas and to highlight the elementary properties of the novel approach. Furthermore, this setting allows for focusing on the core theoretical challenges for truncated variational approximations without unduly complicating the required derivations. In the past, other variational approximations (primarily factored ones) have also been considered for general graphical models or generalized to fully Bayesian settings (e.g. Attias, 2000; Jaakkola, 2000; Beal and Ghahramani, 2003). While the latter would represent a major additional research effort (we focus on discrete latents, for instance), we note that an application of truncated variational EM to graphical models with hierarchical structure or, e.g., for time-series models is straight-forward.

2 Variational Approaches to Expectation Maximization

Before we introduce the novel type of truncated approximations, we first set the ground by introducing variational approximations in general, and by discussing their by far most wide-spread uses in the forms of factored variational EM and Gaussian variational EM. While most of this section reviews well-known results, we will later show that some standard derivations can and have to be generalized for truncated approximations. Furthermore, reviews of other variational approaches facilitates the introduction of terminology and allows to contrast the properties of novel approaches with existing ones.

2.1 The Free-Energy Formulation

Following the framework of Expectation Maximization (EM; Dempster et al., 1977), our aim is to maximize the data likelihood defined by a set of N data points, $\{\vec{y}^{(n)}\}_{n=1,\ldots,N}$, and a generative data model:

$$\Theta^* \;=\; \underset{\Theta}{\operatorname{argmax}}\{L(\Theta)\} \quad \text{with} \quad \mathcal{L}(\Theta) \,=\, \log\Big(p(\vec{y}^{(1)},\ldots,\vec{y}^{(N)}\,|\,\Theta)\Big)\,, \tag{1}$$

where Θ are the parameters of the generative model and where the data points $\vec{y}^{(n)}$ are assumed to be generated independently from a stationary process. The basic EM algorithm was generalized to provide justification for incremental or online versions and, more importantly, to provide the theoretical foundation of variational EM approximations (see, e.g., Hathaway, 1986; Saul et al., 1996; Neal and Hinton, 1998). As it will be of importance further below, we here briefly recapitulate the derivation of the free-energy in a standard form (see, e.g., Saul et al., 1996; Murphy, 2012). For this recall Jensen's inequality, which can for our purposes be denoted as follows: Let f and g be real-valued functions such that $g(\vec{s}), f(\vec{s}) \in \mathbb{R}$ for all $\vec{s} \in \Omega$ and let $q(\vec{s})$ be non-negative and sum to one, i.e., $\sum_{\vec{s}\in\Omega} q(\vec{s}) = 1$. Then for any *concave* function f the following inequality holds:

$$f\big(\sum_{\vec{s}} q(\vec{s})\, g(\vec{s})\big) \;\geq\; \sum_{\vec{s}} q(\vec{s})\, f(g(\vec{s}))\,. \tag{2}$$

Note that $\sum_{\vec{s}}$ in (2) denotes a summation over all states $\vec{s}$, a convention which we will follow also in the remainder of this work.

Any distribution $q(\vec{s})$ on Ω satisfies the conditions for q for Jensen's inequality. If we additionally demand $q(\vec{s})$ to be strictly positive, $q(\vec{s}) > 0$ for all $\vec{s}$, we can apply Jensen's inequality to the data likelihood (1) and obtain by using that the logarithm is concave:

$$\begin{aligned}
\mathcal{L}(\Theta) \;&=\; \sum_{n=1}^{N} \log\big(\sum_{\vec{s}} p(\vec{s},\vec{y}^{(n)}\,|\,\Theta)\big) \\
&=\; \sum_{n=1}^{N} \log\big(\sum_{\vec{s}} q^{(n)}(\vec{s})\,\frac{1}{q^{(n)}(\vec{s})}\,p(\vec{s},\vec{y}^{(n)}\,|\,\Theta)\big) \\
&\geq\; \sum_{n=1}^{N} \sum_{\vec{s}} q^{(n)}(\vec{s})\,\log\big(\frac{1}{q^{(n)}(\vec{s})}\,p(\vec{s},\vec{y}^{(n)}\,|\,\Theta)\big) \\
&=\; \sum_{n=1}^{N} \Big(\sum_{\vec{s}} q^{(n)}(\vec{s})\,\log\big(p(\vec{s},\vec{y}^{(n)}\,|\,\Theta)\big) \;-\; \sum_{\vec{s}} q^{(n)}(\vec{s})\,\log\big(q^{(n)}(\vec{s})\big)\Big)\,,
\end{aligned} \tag{3}$$

where $\sum_{\vec{s}}$ again denotes the sum over all states of $\vec{s}$. The crucial novel entity that emerges in this derivation is a lower bound of the likelihood which is termed the (variational) *free-energy* (compare Saul and Jordan, 1996; Neal and Hinton, 1998; Jordan et al., 1999; MacKay, 2003):

$$\mathcal{L}(\Theta) \;\geq\; \mathcal{F}(q,\Theta) \;:=\; \sum_{n=1}^{N}\Big(\sum_{\vec{s}}\; q^{(n)}(\vec{s})\;\log\big(p(\vec{s},\vec{y}^{(n)}\,|\,\Theta)\big)\Big) + H(q)\,, \tag{4}$$

where $H(q) = -\sum_n \sum_{\vec{s}} q^{(n)}(\vec{s}) \log(q^{(n)}(\vec{s}))$ is a function (the Shannon entropy) that depends on the distribution $q(\vec{s}) = (q^{(1)}(\vec{s}), \ldots, q^{(N)}(\vec{s}))$ but not on the model parameters Θ. Given a data point $\vec{y}^{(n)}$, the distribution $q^{(n)}(\vec{s})$ is called *variational distribution* for reasons that will be discussed below. Also the collection of all distributions $q(\vec{s})$ may be referred to as variational distribution. The name "free-energy" was inherited from statistical physics where approximations to the Helmholtz free-energy of a physical system take a very similar mathematical form (see, e.g., MacKay, 2003). More precisely, while the analog to the Helmholtz free-energy in physics is the data log-likelihood, the analog to variational free-energies in physics is the variational free-energy in (4).

In Machine Learning, the sum over all states of $\vec{s}$ in (4) becomes an integral if the values of $\vec{s}$ are continuous. In the EM scheme $\mathcal{F}(q, \Theta)$ is maximized alternately with respect to q in the E-step (while Θ is kept fixed) and with respect to Θ in the M-step (while q is kept fixed). Using EM it can thus be shown that the EM iterations increase the likelihood or keep it constant. In practical applications EM is found to increase the likelihood to (at least local) likelihood maxima.

A basic result (e.g. Jordan et al., 1999; Neal and Hinton, 1998; MacKay, 2003; Murphy, 2012) is that the difference between the likelihood (1) and the free-energy (4) is given by the Kullback-Leibler divergence between the variational distributions and the posterior distributions:

$$\mathcal{L}(\Theta) - \mathcal{F}(q, \Theta) \;=\; \sum_n D_{\mathrm{KL}}\big(q^{(n)}(\vec{s}), p(\vec{s}\,|\,\vec{y}^{(n)}\,|\,\Theta)\big) \;\geq\; 0, \tag{5}$$

where $D_{\mathrm{KL}}\big(q, p\big)$ denotes the Kullback-Leibler (KL) divergence (see Appendix for the derivation and the definition of the KL-divergence). Eqn. 5 is again analogous to the difference between variational free-energy and Helmholtz free-energy in statistical physics. Using (5) we observe that the free-energy $\mathcal{F}(q, \Theta)$ is maximized if we choose (for all n) the posterior distribution as variational distribution:

$$q^{(n)}(\vec{s}) \;\;:=\;\; p(\vec{s}\,|\,\vec{y}^{(n)}, \Theta) \quad \Rightarrow \quad \mathcal{L}(\Theta) = \mathcal{F}(q, \Theta)\,. \tag{6}$$

While this choice of $q^{(n)}(\vec{s})$ recovers the exact EM algorithm, the computation of posterior distributions represents the crucial computational intractability for most generative data models: typically neither the posteriors nor expectation values w.r.t. them are computationally tractable as they require summations over all possible states of $\vec{s}$.

2.2 The Variational Approximation of EM

Overcoming the problem of computational tractability while maintaining an as good as possible approximation to the true posterior is the main motivation for variational approximations (and other approximation schemes). For variational EM the basic idea is to maximize the free-energy $\mathcal{F}(q, \Theta)$ in (4) instead of directly optimizing the likelihood. If we can find variational distributions q such that

(A) optimization of $\mathcal{F}(q, \Theta)$ is tractable, and

(B) such that the lower bound $\mathcal{F}(q, \Theta)$ becomes as similar (as tight) as possible to $\mathcal{L}(\Theta)$,

a tractable approximate optimization of $\mathcal{L}(\Theta)$ is obtained.

By definition of the free-energy, almost no restrictions are imposed on the choice of the distributions q, such that they can, in principle, be chosen to take any functional form and to be dependent on any set of parameters. To fulfill requirement (A) some choice for a functional form of q has to be made, however. As a consequence, the distributions $q^{(n)}(\vec{s})$ become equipped with additional parameters and these parameters are then optimized to make $\mathcal{F}(q, \Theta)$ as tight as possible (requirement B). The variational distributions are then typically denoted by $q^{(n)}(\vec{s}, \Lambda)$ with the additional parameters Λ being referred to as *variational parameters*. Having chosen a variational distribution, the free-energy is often taken to depend on Λ rather than on the variational distributions themselves, i.e., $\mathcal{F}(\Lambda, \Theta)$.

Without being more specific about the choice of variational distributions and parameters, the iterative procedure to optimize the free-energy may then be abstractly denoted by:

$$\begin{aligned} &\text{Opt 1}: && \Lambda^{\text{new}} = \underset{\Lambda}{\text{argmax}}\mathcal{F}(\Lambda, \Theta) && \text{while holding } \Theta \text{ fixed} \\ &\text{Opt 2}: && \Theta^{\text{new}} = \underset{\Theta}{\text{argmax}}\mathcal{F}(\Lambda^{\text{new}}, \Theta) && \text{while holding } \Lambda^{\text{new}} \text{ fixed} \end{aligned}$$

The two optimization steps are repeated after setting $\Theta = \Theta^{\text{new}}$ until the parameters Θ have sufficiently converged.

The second optimization step (the M-step) is a standard optimization of model parameters while the first optimization (the variational E-step) is an additional procedure that updates the variational parameters. One way to optimize $\mathcal{F}(\Lambda, \Theta)$ is to vary Λ and to choose the Λ corresponding to the largest $\mathcal{F}(\Lambda, \Theta)$ under such a variation. This procedure can often be facilitated based on analytical results. Note that instead of being fully maximized, $\mathcal{F}(\Lambda, \Theta)$ can also just be partially increased in each step. The free-energy would still continuously increase in this case, and the procedure would still be referred to as variational EM.

In principle, many possible choices of variational distributions that fulfill requirements (A) and (B) are conceivable; and any choice would give rise to a variational EM procedure. Still, there are essentially only two standard lines of research on variational EM approaches followed in the literature, one corresponding to the choice of Gaussian variational distributions (*Gaussian Variational EM*) and one corresponding to factored variational distributions (*Factored Variational EM*).

2.3 Gaussian Variational EM (GV-EM)

As one of the most basic distributions is the Gaussian distribution, a natural choice for variational distributions (for continuous latents) is a multi-variate Gaussian:

$$q^{(n)}(\vec{s}) \;:=\; q^{(n)}(\vec{s}; \Lambda) \;=\; q^{(n)}(\vec{s}; \mu^{(n)}, \Sigma^{(n)}) \;=\; \mathcal{N}(\vec{s}; \mu^{(n)}, \Sigma^{(n)})\,, \tag{7}$$

where $\Lambda^{(n)} = (\mu^{(n)}, \Sigma^{(n)})$ and where $\Lambda = (\Lambda^{(1)}, \ldots, \Lambda^{(N)})$ is the set of all variational parameters (one mean and one covariance matrix per data point). The Gaussian variational approximation approximates each posterior distribution by a Gaussian distribution. The optimization of the variational parameters in (7) results in update equations for mean and variance that maximize the free-energy and minimize the KL-divergence between true posteriors and the variational Gaussians. Gaussian distributions are especially well suited for data models with monomodal posteriors. They are popular, for instance, for optimization of sparse linear models (e.g.

Opper and Winther, 2005; Seeger, 2008; Opper and Archambeau, 2009). Gaussians can capture data correlations and jointly optimize mean and variance in the KL-divergence sense. Note that the latter is different from Laplace approximations which first fit the mean and then the covariance matrix (see, e.g., Bishop, 2006) .

2.4 Factored Variational EM (FV-EM)

The second and more predominant line of research on variational EM builds up on the choice of variational distributions that factor over sets of hidden variables. Most commonly the choice is a fully factored distributions:

$$q^{(n)}(\vec{s}) \;\; := \;\; q^{(n)}(\vec{s};\Lambda) \;\; = \;\; \prod_{h=1}^{H} q_h(s_h;\vec{\lambda}_h^{(n)}) \tag{8}$$

where $\vec{\lambda}_h^{(n)}$ are parameters associated with one hidden variable h for one data point n, and where Λ is the collection of all these parameters (all h and n combinations). Using factored distributions in combination with specific choices for the factors $q_h(s_h;\vec{\lambda}_h^{(n)})$ then results in computationally tractable optimizations. For the usual generative models the factors are often chosen to be identical to the prior distributions of the individual hidden variables (e.g. Jordan et al., 1999; Haft et al., 2004).

Generalizations of fully factored approaches can be defined by allowing for dependencies between small sets of variables (doubles, triple etc) resulting in *partially factored approaches* . Such approaches can potentially capture more complex posterior interdependencies, and they are therefore also termed *structured variational approaches* (compare Saul and Jordan, 1996; MacKay, 2003; Bouchard-Côté and Jordan, 2009; Murphy, 2012). The (much more frequently used) fully factored variational approach (8) is also often termed *mean field* approximation (in analogy to statistical physics). Likewise, structured variational approaches are often termed *structured mean-field* to highlight their close relation to fully factored approaches (e.g. Bouchard-Côté and Jordan, 2009; Murphy, 2012). Factored approaches and especially the mean field approach are, by far, the most popular variational approaches. This is also reflected by them being often equated with "variational EM" although a specific factorizing choice was made.

3 Truncated Variational Distributions

The theoretical introduction and the examples of Gaussian and factored variational distributions now provide the ground for the introduction of a novel class of variational distributions. In contrast to the prominent examples of variational EM, we will here neither assume monomodal posteriors (like GV-EM) nor independent factors (like FV-EM) for the definition of variational distributions. Instead, we choose the variational distributions to be proportional to the exact posterior distributions themselves but only on a subset of hidden states $\vec{s}$. The variational distri-

bution is then defined as a truncated distribution of the following form:

$$q^{(n)}(\vec{s}) := q^{(n)}(\vec{s};\mathcal{K},\Theta) = \frac{p(\vec{s}\,|\,\vec{y}^{\,(n)},\Theta)}{\sum\limits_{\vec{s}\,'\in\mathcal{K}^{(n)}} p(\vec{s}\,'\,|\,\vec{y}^{\,(n)},\Theta)}\,\delta(\vec{s}\in\mathcal{K}^{(n)}) = \frac{p(\vec{s},\vec{y}^{\,(n)}\,|\,\Theta)}{\sum\limits_{\vec{s}\,'\in\mathcal{K}^{(n)}} p(\vec{s}\,',\vec{y}^{\,(n)}\,|\,\Theta)}\,\delta(\vec{s}\in\mathcal{K}^{(n)}),\quad (9)$$

where $\delta(\vec{s}\in\mathcal{K}^{(n)})$ is an indicator function, i.e., $\delta(\vec{s}\in\mathcal{K}^{(n)}) = 1$ if $\vec{s}\in\mathcal{K}^{(n)}$ and zero otherwise. As was the case for Gaussian and factored variational distributions, we have one set of variational parameters per data point. We therefore define $\mathcal{K} = (\mathcal{K}^{(1)},\ldots,\mathcal{K}^{(N)})$ to be the collection of all sets $\mathcal{K}^{(n)}$ (see Eqn. 9). Truncated distributions have been suggested previously and have successfully be applied to a number of elementary as well as more complex generative data model descriptions (Lücke and Eggert, 2010; Dai et al., 2013; Dai and Lücke, 2014; Sheikh et al., 2014; Sheikh and Lücke, 2016). Instead of using a variational optimization of approximation parameters similar to GV-EM or TV-EM, truncated approximations have, so far, used preselection mechanisms to reduce the number of states evaluated for a truncated approximation (Lücke and Eggert, 2010). No variational optimization loop was used and it was unclear how such an optimization could be defined (we will address this central problem in Sec. 4).

By considering Eqn. 9, it can instantly be observed that $q^{(n)}(\vec{s})$ becomes equal to the exact posterior $p(\vec{s}\,|\,\vec{y}^{\,(n)},\Theta)$ if $\mathcal{K}^{(n)}$ contains all possible states of $\vec{s}$. In general, computing all latent states is computationally prohibitive such that $\mathcal{K}^{(n)}$ has to be a sufficiently small subset, and $q^{(n)}(\vec{s})$ becomes an approximation to the posterior. The larger the posterior mass that can be captured by the states in $\mathcal{K}^{(n)}$, the better is the approximation. Consequently, if the posterior mass is concentrated in small volumes of the state space of $\vec{s}$, approximation (9) can become very accurate for small $\mathcal{K}^{(n)}$.

For the truncated variational approximation to become practically applicable for approximate EM algorithms, expectation values w.r.t. the variational distribution $q^{(n)}(\vec{s};\mathcal{K},\Theta)$ have to be efficiently computable. By using (9) as a distribution, the expectation value of a function $g(\vec{s})$ over latents $\vec{s}$ is given by:

$$\langle g(\vec{s})\rangle_{q^{(n)}} \;=\; \langle g(\vec{s})\rangle_{q^{(n)}(\vec{s};\mathcal{K},\Theta)} \;=\; \frac{\sum\limits_{\vec{s}\in\mathcal{K}^{(n)}} p(\vec{s},\vec{y}^{\,(n)}\,|\,\Theta)\;g(\vec{s})}{\sum\limits_{\vec{s}\,'\in\mathcal{K}^{(n)}} p(\vec{s}\,',\vec{y}^{\,(n)}\,|\,\Theta)}\,. \quad (10)$$

As the joint $p(\vec{s},\vec{y}^{\,(n)}\,|\,\Theta)$, e.g., of a probabilistic generative model is efficiently computable, expectation values (10) are computationally tractable if $\mathcal{K}^{(n)}$ contains a sufficiently low number of states.

Having defined the variational distribution of (9), we can now seek to derive the corresponding free-energy. However, by considering the standard derivation (3) showing that the free-energy is a lower bound of the likelihood, recall that we required the values of $q^{(n)}(\vec{s})$ to be strictly positive. To consider truncated distributions (9) within the free-energy framework, we therefore first have to generalize the free-energy formalism to be applicable to any distribution in the latent space.

As a fist step, we therefore show that the following proposition holds.

Proposition 1
Let $q^{(n)}(\vec{s})$ be variational distributions defined on a set of states Ω (with values not necessarily greater zero), then a free-energy function $\mathcal{F}(q,\Theta)$ exists and is given by:

$$\mathcal{F}(q,\Theta) := \sum_{n=1}^{N}\Big(\sum_{\vec{s}} q^{(n)}(\vec{s})\ \log\big(p(\vec{s},\vec{y}^{(n)}\,|\,\Theta)\big)\Big) + H(q)\,. \tag{11}$$

Proof
First let us (as in the Proposition) drop the notation of variational parameters to increase readability. Then observe that the standard derivation (3) shows that the proposition is true if $q^{(n)}(\vec{s}) > 0$ for all $\vec{s}\in\Omega$ and all n. To show that (11) is true for any distribution, consider the case of variational distributions $q^{(n)}(\vec{s})$ for which $q^{(n)}(\vec{s}) > 0$ is true only in proper subsets $\mathcal{K}^{(n)}$ of the state space Ω and equal to zero for all $\vec{s}\notin\mathcal{K}^{(n)}$. For a given n a distribution is either $q^{(n)}(\vec{s}) > 0$ or there exists a proper subset $\mathcal{K}^{(n)}$. In the latter case it is trivial to define such a $\mathcal{K}^{(n)}$, and as $q^{(n)}(\vec{s})$ is a distribution (i.e., sums to one), this implies that $\mathcal{K}^{(n)}$ is not empty. For the main part of the proof we will now assume that for all n there exists a proper subset $\mathcal{K}^{(n)}$ of Ω. A general q may consist of distributions $q^{(n)}(\vec{s})$ that are strictly positive for some n and have proper subsets $\mathcal{K}^{(n)}$ for the other n. However, addressing these mixed cases will turn out to be straight-forward, and we will come back to this general case at the end of the proof.

Given a distribution $q^{(n)}(\vec{s})$ with set $\mathcal{K}^{(n)}$, let us define an auxiliary function $\tilde{q}^{(n)}(\vec{s})$ as follows:

$$\tilde{q}^{(n)}(\vec{s}) = \begin{cases} q^{(n)}(\vec{s}) - \epsilon_n^- & \text{for all } \vec{s}\in\mathcal{K}^{(n)} \\ q^{(n)}(\vec{s}) + \epsilon_n^+ & \text{for all } \vec{s}\notin\mathcal{K}^{(n)} \end{cases} \tag{12}$$

Let ϵ_n^- be greater zero but smaller than any value of $q^{(n)}(\vec{s})$ in $\mathcal{K}^{(n)}$, i.e.,

$$0 < \epsilon_n^- < \min_{\vec{s}\in\mathcal{K}^{(n)}}\{q^{(n)}(\vec{s})\} \qquad \text{and let} \qquad \epsilon_n^+ := \frac{|\mathcal{K}^{(n)}|}{|\Omega| - |\mathcal{K}^{(n)}|}\,\epsilon_n^-\,, \tag{13}$$

where $|\mathcal{K}|$ is the number of states in $\mathcal{K}$ and where $|\Omega|$ is the number of all possible states, e.g., 2^H for H binary variables. As we have defined $\mathcal{K}^{(n)}$ to contain only values $q^{(n)}(\vec{s}) > 0$, the minimum in $\mathcal{K}^{(n)}$ is greater zero, i.e., we can always find an ϵ_n^- satisfying (13). Consequently, ϵ_n^+ is also greater zero and finite as we demanded $\mathcal{K}^{(n)}$ to be a proper subsets of Ω ($|\mathcal{K}^{(n)}| < |\Omega|$).

With definitions (13) observe that $\tilde{q}^{(n)}(\vec{s}) > 0$ for all $\vec{s}$ and that $\sum_{\vec{s}}\tilde{q}(\vec{s}) = 1$, i.e., $\tilde{q}^{(n)}(\vec{s})$ is distribution on Ω which satisfies (other than $q^{(n)}(\vec{s})$) the requirement for original derivation of the variational free-energy (3). We can therefore use the free-energy definition for $\tilde{q}(\vec{s}) = \big(\tilde{q}^{(1)}(\vec{s}),\ldots,\tilde{q}^{(N)}(\vec{s})\big)$ and then consider the limit of small ϵ_n^-. For this, we insert the definition of

$\tilde{q}^{(n)}(\vec{s})$ into the free-energy and find:

$$
\begin{aligned}
&\mathcal{L}(\Theta) \geq \mathcal{F}(\tilde{q},\Theta)\\
&= \sum_n \Big(\sum_{\vec{s}} \tilde{q}^{(n)}(\vec{s}) \log\big(p(\vec{s},\vec{y}^{\,(n)}\,|\,\Theta)\big) - \sum_{\vec{s}} \tilde{q}^{(n)}(\vec{s}) \log\big(\tilde{q}^{(n)}(\vec{s})\big)\Big)\\
&= \sum_n \Big(\sum_{\vec{s}\in\mathcal{K}^{(n)}} \big(q^{(n)}(\vec{s})-\epsilon_n^-\big) \log\big(p(\vec{s},\vec{y}^{\,(n)}\,|\,\Theta)\big) - \sum_{\vec{s}\in\mathcal{K}^{(n)}} \big(q^{(n)}(\vec{s})-\epsilon_n^-\big) \log\big(q^{(n)}(\vec{s})-\epsilon_n^-\big)\\
&\quad + \sum_{\vec{s}\notin\mathcal{K}^{(n)}} \big(q^{(n)}(\vec{s})+\epsilon_n^+\big) \log\big(p(\vec{s},\vec{y}^{\,(n)}\,|\,\Theta)\big) - \sum_{\vec{s}\notin\mathcal{K}^{(n)}} \big(q^{(n)}(\vec{s})+\epsilon_n^+\big) \log\big(q^{(n)}(\vec{s})+\epsilon_n^+\big)\Big)\\
&= \sum_n \Big(\sum_{\vec{s}} q^{(n)}(\vec{s}) \log\big(p(\vec{s},\vec{y}^{\,(n)}\,|\,\Theta)\big) - \epsilon_n^- \sum_{\vec{s}\in\mathcal{K}^{(n)}} \log\big(p(\vec{s},\vec{y}^{\,(n)}\,|\,\Theta)\big) + \epsilon_n^+ \sum_{\vec{s}\notin\mathcal{K}^{(n)}} \log\big(p(\vec{s},\vec{y}^{\,(n)}\,|\,\Theta)\big)\\
&\quad - \sum_{\vec{s}\in\mathcal{K}^{(n)}} \big(q^{(n)}(\vec{s})-\epsilon_n^-\big) \log\big(q^{(n)}(\vec{s})-\epsilon_n^-\big) - \epsilon_n^+ \log(\epsilon_n^+) \sum_{\vec{s}\notin\mathcal{K}^{(n)}} 1 \quad \Big)\\
&= \sum_n \Big(\sum_{\vec{s}} q^{(n)}(\vec{s}) \log\big(p(\vec{s},\vec{y}^{\,(n)}\,|\,\Theta)\big) - \epsilon_n^- \sum_{\vec{s}\in\mathcal{K}^{(n)}} \log\big(p(\vec{s},\vec{y}^{\,(n)}\,|\,\Theta)\big) + \epsilon_n^- \frac{|\mathcal{K}^{(n)}|}{|\Omega|-|\mathcal{K}^{(n)}|} \sum_{\vec{s}\notin\mathcal{K}^{(n)}} \log\big(p(\vec{s},\vec{y}^{\,(n)}\,|\,\Theta)\big)\\
&\quad - \sum_{\vec{s}\in\mathcal{K}^{(n)}} \big(q^{(n)}(\vec{s})-\epsilon_n^-\big) \log\big(q^{(n)}(\vec{s})-\epsilon_n^-\big) - \epsilon_n^- |\mathcal{K}^{(n)}| \log\big(\frac{|\mathcal{K}^{(n)}|}{|\Omega|-|\mathcal{K}^{(n)}|}\big) - |\mathcal{K}^{(n)}|\,\epsilon_n^- \log(\epsilon_n^-)\Big)\\
&= \sum_n \Big(\sum_{\vec{s}} q^{(n)}(\vec{s}) \log\big(p(\vec{s},\vec{y}^{\,(n)}\,|\,\Theta)\big) - \sum_{\vec{s}\in\mathcal{K}^{(n)}} \big(q^{(n)}(\vec{s})-\epsilon_n^-\big) \log\big(q^{(n)}(\vec{s})-\epsilon_n^-\big) - |\mathcal{K}^{(n)}|\,\epsilon_n^- \log(\epsilon_n^-)\\
&\quad - \epsilon_n^- \sum_{\vec{s}\in\mathcal{K}^{(n)}} \log\big(p(\vec{s},\vec{y}^{\,(n)}\,|\,\Theta)\big) + \epsilon_n^- \frac{|\mathcal{K}^{(n)}|}{|\Omega|-|\mathcal{K}^{(n)}|} \sum_{\vec{s}\notin\mathcal{K}^{(n)}} \log\big(p(\vec{s},\vec{y}^{\,(n)}\,|\,\Theta)\big) - \epsilon_n^- |\mathcal{K}^{(n)}| \log\big(\frac{|\mathcal{K}^{(n)}|}{|\Omega|-|\mathcal{K}^{(n)}|}\big)\Big)
\end{aligned}
$$

This expression applies for any ϵ_n^- satisfying (13) and therefore also if we replace:

$$\epsilon_n^- \;=\; \epsilon \;<\; \min_{n'=1,\ldots,N}\{\epsilon_{n'}^-\}, \tag{14}$$

such that we obtain:

$$
\begin{aligned}
\mathcal{F}(\tilde{q},\Theta) = \sum_n \Big(\sum_{\vec{s}} q^{(n)}(\vec{s}) \log\big(p(\vec{s},\vec{y}^{\,(n)}\,|\,\Theta)\big) - \sum_{\vec{s}\in\mathcal{K}^{(n)}} \big(q^{(n)}(\vec{s})-\epsilon\big) \log\big(q^{(n)}(\vec{s})-\epsilon\big) - |\mathcal{K}^{(n)}|\,\epsilon \log(\epsilon)\\
- \epsilon \sum_{\vec{s}\in\mathcal{K}^{(n)}} \log\big(p(\vec{s},\vec{y}^{\,(n)}\,|\,\Theta)\big) + \epsilon \frac{|\mathcal{K}^{(n)}|}{|\Omega|-|\mathcal{K}^{(n)}|} \sum_{\vec{s}\notin\mathcal{K}^{(n)}} \log\big(p(\vec{s},\vec{y}^{\,(n)}\,|\,\Theta)\big) - \epsilon\,|\mathcal{K}^{(n)}| \log\big(\frac{|\mathcal{K}^{(n)}|}{|\Omega|-|\mathcal{K}^{(n)}|}\big)\Big)
\end{aligned}
\tag{15}
$$

Let us now consider infinitesimally small $\epsilon > 0$, i.e., let us consider the limit when $\epsilon \to 0$. First, observe that in this case all summands of the last line of $\mathcal{F}(\tilde{q},\Theta)$ in (15) trivially converge to zero. Then observe that the second summand of $\mathcal{F}(\tilde{q},\Theta)$ converges to

$$\lim_{\epsilon\to 0} \Big(\sum_{\vec{s}\in\mathcal{K}^{(n)}} \big(q^{(n)}(\vec{s}) - \epsilon\big) \log\big(q^{(n)}(\vec{s}) - \epsilon\big)\Big) = \sum_{\vec{s}\in\mathcal{K}^{(n)}} q^{(n)}(\vec{s}) \log\big(q^{(n)}(\vec{s})\big) \tag{16}$$

as $q^{(n)}(\vec{s})$ is finite and greater zero for all $\vec{s} \in \mathcal{K}^{(n)}$. Finally, the third summand in (15) can be observed to converge (following l'Hôpital) to zero:

$$\lim_{\epsilon\to 0} \big(\epsilon \log(\epsilon)\big) \;=\; \lim_{\epsilon\to 0} \big(\frac{\log(\epsilon)}{\frac{1}{\epsilon}}\big) \;=\; \lim_{\epsilon\to 0} \big(-\frac{\frac{1}{\epsilon}}{\frac{1}{\epsilon^2}}\big) \;=\; -\lim_{\epsilon\to 0} \big(\epsilon\big) \;=\; 0\,. \tag{17}$$

Thus, the limit $\epsilon \to 0$ exists and is given by:

$$\mathcal{F}(q,\Theta) = \lim_{\epsilon\to 0}\mathcal{F}(\tilde{q},\Theta) = \sum_n\sum_{\vec{s}} q^{(n)}(\vec{s})\,\log\left(p(\vec{s},\vec{y}^{\,(n)}\,|\,\Theta)\right) - \sum_n\sum_{\vec{s}} q^{(n)}(\vec{s})\,\log\left(q^{(n)}(\vec{s})\right) \tag{18}$$

where we have used because of (17) the convention $q^{(n)}(\vec{s})\,\log\left(q^{(n)}(\vec{s})\right) = 0$ for all $q^{(n)}(\vec{s}) = 0$. We will use this convention (which is also commonly used for the KL-divergence) throughout the paper.

Eqn. 18 shows that Proposition 1 holds for any q with distributions $q^{(n)}(\vec{s})$ with proper subsets $\mathcal{K}^{(n)}$ of Ω. To finally show that Proposition 1 hold for general (discrete) q, consider the mixed case that q contains strictly positive distributions $q^{(n)}$ and distributions $q^{(n)}$ with exact zeros. Let us define the set $J \subseteq \{1,\ldots,N\}$ to contain those n with $q^{(n)}$ being strictly positive, and the complement $\overline{J} \subseteq \{1,\ldots,N\}$ to contain those n with $q^{(n)}$ that are not strictly positive. Using J and $\overline{J}$ we then define an auxiliary function $\tilde{q}$ by using $\tilde{q}^{(n)}(\vec{s})$ of Eqn. 12 only for $n \in \overline{J}$ and by setting $\tilde{q}^{(n)}(\vec{s}) = q^{(n)}(\vec{s})$ for all $\vec{s}$ for all $n \in J$. The functions $\tilde{q}^{(n)}$ are then again strictly positive distributions on Ω for all n and we can again apply the standard result (3):

$$\begin{aligned}
\mathcal{L}(\Theta) \geq \mathcal{F}(\tilde{q},\Theta) &= \sum_n\Big(\sum_{\vec{s}} \tilde{q}^{(n)}(\vec{s})\,\log\left(p(\vec{s},\vec{y}^{\,(n)}\,|\,\Theta)\right) - \sum_{\vec{s}} \tilde{q}^{(n)}(\vec{s})\,\log\left(\tilde{q}^{(n)}(\vec{s})\right)\Big)\\
&= \sum_{n\in J}\Big(\sum_{\vec{s}} q^{(n)}(\vec{s})\,\log\left(p(\vec{s},\vec{y}^{\,(n)}\,|\,\Theta)\right) - \sum_{\vec{s}} q^{(n)}(\vec{s})\,\log\left(q^{(n)}(\vec{s})\right)\Big)\\
&\quad + \sum_{n\in\overline{J}}\Big(\sum_{\vec{s}} \tilde{q}^{(n)}(\vec{s})\,\log\left(p(\vec{s},\vec{y}^{\,(n)}\,|\,\Theta)\right) - \sum_{\vec{s}} \tilde{q}^{(n)}(\vec{s})\,\log\left(\tilde{q}^{(n)}(\vec{s})\right)\Big)
\end{aligned}$$

Now we apply the same arguments as above but only to the sum over all $n \in \overline{J}$, for which everything remains as for the derivation above. Eqn. 18 therefore applies if we only consider sums over $\overline{J}$. Thus in the limit $\epsilon \to 0$ we obtain in the mixed case:

$$\begin{aligned}
\mathcal{F}(q,\Theta) &= \lim_{\epsilon\to 0}\left(\mathcal{F}(\tilde{q},\Theta)\right)\\
&= \sum_{n\in J}\Big(\sum_{\vec{s}} q^{(n)}(\vec{s})\,\log\left(p(\vec{s},\vec{y}^{\,(n)}\,|\,\Theta)\right) - \sum_{\vec{s}} q^{(n)}(\vec{s})\,\log\left(q^{(n)}(\vec{s})\right)\Big)\\
&\quad + \lim_{\epsilon\to 0}\Big(\sum_{n\in\overline{J}}\Big(\sum_{\vec{s}} \tilde{q}^{(n)}(\vec{s})\,\log\left(p(\vec{s},\vec{y}^{\,(n)}\,|\,\Theta)\right) - \sum_{\vec{s}} \tilde{q}^{(n)}(\vec{s})\,\log\left(\tilde{q}^{(n)}(\vec{s})\right)\Big)\Big)\\
&= \sum_n\Big(\sum_{\vec{s}} q^{(n)}(\vec{s})\,\log\left(p(\vec{s},\vec{y}^{\,(n)}\,|\,\Theta)\right) - \sum_{\vec{s}} q^{(n)}(\vec{s})\,\log\left(q^{(n)}(\vec{s})\right)\Big)
\end{aligned}$$

which finally shows that Proposition 1 holds for any q with any distributions $q^{(n)}$ on Ω.
□

As the free-energy (11) for general q was obtained as a limit, we have to show that it remains a lower bound of $\mathcal{L}(\Theta)$ also in this limit. Using the KL-divergence result (5), this can however be shown using a similar approach as for the proof above.

Proposition 2
Let $q^{(n)}(\vec{s})$ be variational distributions over a set of states Ω (with values not necessarily greater zero), then the corresponding free-energy (11) is a lower bound of the likelihood, and the difference

between likelihood and free-energy is given by the sum over KL-divergences:

$$\mathcal{L}(\Theta) - \mathcal{F}(q,\Theta) \;=\; \sum_n D_{\mathrm{KL}}\big(q^{(n)}(\vec{s}), p(\vec{s}\,|\,\vec{y}^{\,(n)}\,|\,\Theta)\big) \;\geq\; 0. \tag{19}$$

Proof
Given distributions $q^{(n)}(\vec{s})$ let us consider the same auxiliary distributions $\tilde{q}^{(n)}(\vec{s})$ as for the proof of proposition 1. As $\tilde{q}^{(n)}(\vec{s})$ are strictly positive distributions, it applies according to the standard derivation (see Appendix):

$$\mathcal{L}(\Theta) - \mathcal{F}(\tilde{q},\Theta) \;=\; \sum_n D_{\mathrm{KL}}\big(\tilde{q}^{(n)}(\vec{s}), p(\vec{s}\,|\,\vec{y}^{\,(n)}\,|\,\Theta)\big) \;\geq\; 0, \tag{20}$$

for any set of ϵ_n^- satisfying (13). If we set all ϵ_n^- equal to a sequence $\epsilon_k = 1/k$, we know that there exist a finite K such that for all $k > K$ the ϵ_n^- fulfill condition (13). If we now take the distribution $\tilde{q}_k(\vec{s}) = (\tilde{q}_k^{(1)}, \ldots, \tilde{q}_k^{(N)})$ to be defined by choosing $\epsilon_n^- = \epsilon_k = 1/k$ for all n, we know that for all $k > K$

$$D_k \;=\; \mathcal{L}(\Theta) - \mathcal{F}(\tilde{q}_k,\Theta) \;=\; \sum_n D_{\mathrm{KL}}\big(\tilde{q}_k^{(n)}(\vec{s}), p(\vec{s}\,|\,\vec{y}^{\,(n)}\,|\,\Theta)\big) \;\geq\; 0. \tag{21}$$

The sequence D_k is hence a sequence in the interval $[0,\infty)$. As the limit $\lim_{k\to\infty}\mathcal{F}(\tilde{q}_k,\Theta)$ is finite, D_k converges to a finite value within $[0,\infty)$ (which is left-closed). If q contains some strictly positive $q^{(n)}$, we only use $\epsilon_n^- = \epsilon_k = 1/k$ for all $n \in \overline{J}$. Finally, by using Proposition 1, the limit $\lim_{k\to\infty} D_k$ is given by:

$$\begin{aligned}
0 \leq \lim_{k\to\infty} D_k \;&=\; \mathcal{L}(\Theta) - \lim_{k\to\infty}\mathcal{F}(\tilde{q}_k,\Theta)\\
&=\; \mathcal{L}(\Theta) \;-\; \mathcal{F}(q,\Theta)\\
&=\; \mathcal{L}(\Theta) \;-\; \sum_n\sum_{\vec{s}} q^{(n)}(\vec{s})\,\log\big(p(\vec{s},\vec{y}^{\,(n)}\,|\,\Theta)\big) \;+\; \sum_n\sum_{\vec{s}} q^{(n)}(\vec{s})\,\log\big(q^{(n)}(\vec{s})\big)\\
&=\; \mathcal{L}(\Theta) - \sum_n \underbrace{\sum_{\vec{s}} q^{(n)}(\vec{s})}_{1}\,\log\big(p(\vec{y}^{\,(n)}\,|\,\Theta)\big) - \sum_n\sum_{\vec{s}} q^{(n)}(\vec{s})\,\log\big(p(\vec{s}\,|\,\vec{y},\Theta)\big) + \sum_n\sum_{\vec{s}} q^{(n)}(\vec{s})\,\log\big(q^{(n)}(\vec{s})\big)\\
&=\; -\sum_n\Big(\sum_{\vec{s}} q^{(n)}(\vec{s})\,\log\big(p(\vec{s}\,|\,\vec{y},\Theta)\big) \;-\; \sum_n\sum_{\vec{s}} q^{(n)}(\vec{s})\,\log\big(q^{(n)}(\vec{s})\big)\Big)\\
&=\; \sum_n D_{\mathrm{KL}}\big(q^{(n)}(\vec{s}), p(\vec{s}\,|\,\vec{y}^{\,(n)}\,|\,\Theta)\big)\,,
\end{aligned}$$

where the last part follows the lines of the standard derivation for the difference $\mathcal{L}(\Theta) - \mathcal{F}(q,\Theta)$. Note that we again used the convention $q^{(n)}(\vec{s})\,\log\big(q^{(n)}(\vec{s})\big) = 0$ for all $q^{(n)}(\vec{s}) = 0$.
□

Taken together, Propositions 1 and 2 mean that we can drop the constraint to strictly positive variational distributions. As a consequence, we can use the (generalized) free-energy framework also for the truncated variational distributions in (9). The variational free-energy associated with a truncated distribution will be referred to as *truncated variational free-energy* or simply *truncated free-energy*.

We can now insert the specific truncated distributions (9) into the free-energy (11). As for non-variational EM (6) with exact posterior as variational distributions, $q^{(n)}(\vec{s}) = p(\vec{s}\,|\,\vec{y}^{\,(n)}, \hat{\Theta})$, we will distinguish between parameters $\hat{\Theta}$ of the variational distribution and the parameters Θ of the generative data model. Like for the M-step of exact EM, this allows for taking derivatives of the log-joint $\log(p(\vec{s}, \vec{y}^{\,(n)}\,|\,\Theta))$ while the variational distributions can be treated as constant (we will come back to this point further below). Inserting the truncated distributions $q^{(n)}(\vec{s}; \mathcal{K}, \hat{\Theta})$ into the free-energy (11) then yields:

$$\hat{\mathcal{F}}(\mathcal{K}, \hat{\Theta}, \Theta) \;=\; \sum_{n=1}^{N} \left[\sum_{\vec{s}} \; q^{(n)}(\vec{s}; \mathcal{K}, \hat{\Theta}) \;\log\left(p(\vec{s}, \vec{y}^{\,(n)}\,|\,\Theta)\right) \right] + H(q(\vec{s}; \mathcal{K}, \hat{\Theta}))\,, \tag{22}$$

The free-energy now depends on three sets of parameters, where the Shannon-entropy $H(q(\vec{s}; \mathcal{K}, \hat{\Theta}))$ is independent of the parameters Θ.

4 Optimization of Truncated Variational Free-Energies

The variational parameters of $\mathcal{K}$ and $\hat{\Theta}$ that we aim at optimizing are different from the typical variational parameters, e.g., of a factored variational approach or a Gaussian approximation. $\mathcal{K} = (\mathcal{K}^{(1)}, \ldots, \mathcal{K}^{(N)})$ are sets of discrete points in latent space; and the parameters $\hat{\Theta}$ are of the same type as those of the generative model but with potentially different values. Still, the distribution $q^{(n)}(\vec{s}; \mathcal{K}, \hat{\Theta})$ can be used within the variational free-energy framework as we have shown above.

Following the free-energy approach we now aim at optimizing the free-energy (22) instead of directly optimizing the likelihood. We will do so by optimizing $\hat{\mathcal{F}}(\mathcal{K}, \hat{\Theta}, \Theta)$ step-by-step w.r.t. its three sets of parameters: %beginarraylclll

$$\begin{array}{lllll}
\text{Opt 1:} & \hat{\Theta}^{\text{new}} & = \underset{\hat{\Theta}}{\text{argmax}} \hat{\mathcal{F}}(\mathcal{K}, \hat{\Theta}, \Theta) & \text{while holding } \mathcal{K} \text{ and } \Theta \text{ fixed} \\
\text{Opt 2:} & \mathcal{K}^{\text{new}} & = \underset{\mathcal{K}}{\text{argmax}} \hat{\mathcal{F}}(\mathcal{K}, \hat{\Theta}^{\text{new}}, \Theta) & \text{while holding } \hat{\Theta}^{\text{new}} \text{ and } \Theta \text{ fixed} \\
\text{Opt 3:} & \Theta^{\text{new}} & = \underset{\Theta}{\text{argmax}} \hat{\mathcal{F}}(\mathcal{K}^{\text{new}}, \hat{\Theta}^{\text{new}}, \Theta) & \text{while holding } \mathcal{K}^{\text{new}} \text{ and } \hat{\Theta}^{\text{new}} \text{ fixed}
\end{array} \tag{23}$$

After setting $\mathcal{K} = \mathcal{K}^{\text{new}}$ and $\Theta = \Theta^{\text{new}}$ we start with the first optimization step again and obtain an iterative procedure. We refer to one iteration step as a truncated variational EM iteration (TV-EM iteration). The first two optimization steps maximize the variational parameters $\hat{\Theta}$ and $\mathcal{K}$, which corresponds to the truncated variational E-step. The third optimization updates the model parameters is the M-step. By definition, a TV-EM iteration never decreases the free-energy $\hat{\mathcal{F}}(\mathcal{K}, \hat{\Theta}, \Theta)$. Also partial optimization, i.e., improving $\hat{\mathcal{F}}(\mathcal{K}, \hat{\Theta}, \Theta)$ instead of maximizing it, in any optimization step does never decrease the free-energy.

The optimization steps (23) of TV-EM are formal definitions. In order to be applicable, a more concrete procedure for each of the three optimization steps is required.

4.1 The Truncated Free-Energy

Instead of investigating and applying the three optimization steps individually, we will carefully analyze each optimization in a theoretically grounded way. For this purpose, let us first introduce

a free-energy defined by setting the values of the variational parameters $\hat{\Theta}$ equal to the model parameters Θ:

$$\mathcal{F}(\mathcal{K},\Theta) := \hat{\mathcal{F}}(\mathcal{K},\Theta,\Theta) = \sum_{n=1}^{N}\Big[\sum_{\vec{s}} q^{(n)}(\vec{s};\mathcal{K},\Theta)\ \log\Big(p(\vec{s},\vec{y}^{\,(n)}\,|\,\Theta)\Big)\Big] + H(q^{(n)}(\vec{s};\mathcal{K},\Theta))\,. \tag{24}$$

Given the definition of the truncated variational distribution $q^{(n)}(\vec{s};\mathcal{K},\Theta)$, it can then be shown that the simplified free-energy $\mathcal{F}(\mathcal{K},\Theta)$ can be expressed in very compact form:

Proposition 3
Given a generative model defined by the joint distribution $p(\vec{s},\vec{y}\,|\,\Theta)$, let $\mathcal{L}(\Theta)$ be the likelihood in Eqn. 1. If $\mathcal{F}(\mathcal{K},\Theta)$ is the free-energy define by Eqn. 24 then it follows that

$$\mathcal{F}(\mathcal{K},\Theta) \;=\; \sum_{n=1}^{N}\log\big(\sum_{\vec{s}\in\mathcal{K}^{(n)}} p(\vec{s},\vec{y}^{\,(n)}\,|\,\Theta)\big). \tag{25}$$

Proof
Following Proposition 1 and 2, $\mathcal{F}(\mathcal{K},\Theta)$ is a lower bound of $\mathcal{L}(\Theta)$, which satisfies:

$$\mathcal{L}(\Theta) - \mathcal{F}(\mathcal{K},\Theta) = \sum_{n} D_{\mathrm{KL}}\big(q^{(n)}(\vec{s};\mathcal{K},\Theta), p(\vec{s}\,|\,\vec{y}^{\,(n)},\Theta)\big) \tag{26}$$

For notational purposes let us introduce the normalizer $Z^{(n)} = \sum_{\vec{s}\in\mathcal{K}^{(n)}} p(\vec{s}\,|\,\vec{y}^{\,(n)},\Theta)$ such that:

$$q^{(n)}(\vec{s};\mathcal{K},\Theta) \;=\; \frac{1}{Z^{(n)}}\,p(\vec{s}\,|\,\vec{y}^{\,(n)}\,|\,\Theta)\,\delta(\vec{s}\in\mathcal{K}^{(n)}) \tag{27}$$

From the above it follows that:

$$\begin{aligned}
&\mathcal{F}(\mathcal{K},\Theta)\\
&\quad= \mathcal{L}(\Theta) - \sum_{n} D_{\mathrm{KL}}\big(q^{(n)}(\vec{s};\mathcal{K},\Theta), p(\vec{s}\,|\,\vec{y}^{\,(n)}\,|\,\Theta)\big)\\
&\quad= \sum_{n}\log(p(\vec{y}^{\,(n)}\,|\,\Theta)) + \sum_{n}\sum_{\vec{s}} q^{(n)}(\vec{s};\mathcal{K},\Theta)\log\big(\frac{p(\vec{s}\,|\,\vec{y}^{\,(n)}\,|\,\Theta)}{q^{(n)}(\vec{s};\mathcal{K},\Theta)}\big)\\
&\quad= \sum_{n}\log(p(\vec{y}^{\,(n)}\,|\,\Theta)) + \sum_{n}\sum_{\vec{s}} \frac{1}{Z^{(n)}}p(\vec{s}\,|\,\vec{y}^{\,(n)}\,|\,\Theta)\,\delta(\vec{s}\in\mathcal{K}^{(n)})\log\big(\frac{p(\vec{s}\,|\,\vec{y}^{\,(n)}\,|\,\Theta)}{\frac{1}{Z^{(n)}}p(\vec{s}\,|\,\vec{y}^{\,(n)}\,|\,\Theta)\,\delta(\vec{s}\in\mathcal{K}^{(n)}))}\big)\\
&\quad= \sum_{n}\log(p(\vec{y}^{\,(n)}\,|\,\Theta)) + \sum_{n}\sum_{\vec{s}\in\mathcal{K}^{(n)}} \frac{1}{Z^{(n)}}\,p(\vec{s}\,|\,\vec{y}^{\,(n)}\,|\,\Theta)\log\big(\frac{p(\vec{s}\,|\,\vec{y}^{\,(n)}\,|\,\Theta)}{\frac{1}{Z^{(n)}}p(\vec{s}\,|\,\vec{y}^{\,(n)}\,|\,\Theta)}\big)\\
&\quad= \sum_{n}\log(p(\vec{y}^{\,(n)}\,|\,\Theta)) + \sum_{n}\sum_{\vec{s}\in\mathcal{K}^{(n)}} \frac{1}{Z^{(n)}}\,p(\vec{s}\,|\,\vec{y}^{\,(n)}\,|\,\Theta)\log\big(Z^{(n)}\big).
\end{aligned} \tag{28}$$

Again we used the convention $q^{(n)}(\vec{s})\,\log\big(q^{(n)}(\vec{s})\big) = 0$ for all $q^{(n)}(\vec{s}) = 0$. Observing the interme-

diate result above, note that $Z^{(n)}$ is independent of $\vec{s}$. We can therefore continue as follows:

$$\begin{aligned}
\mathcal{F}(\mathcal{K},\Theta) &= \sum_n \log(p(\vec{y}^{\,(n)}\,|\,\Theta)) + \sum_n \log\left(Z^{(n)}\right)\frac{\sum_{\vec{s}\in\mathcal{K}^{(n)}} p(\vec{s}\,|\,\vec{y}^{\,(n)},\Theta)}{Z^{(n)}} \\
&= \sum_n \log(p(\vec{y}^{\,(n)}\,|\,\Theta)) + \sum_n \log\left(Z^{(n)}\right) \\
&= \sum_n \log(p(\vec{y}^{\,(n)}\,|\,\Theta)) + \sum_n \log\Big(\sum_{\vec{s}\in\mathcal{K}^{(n)}} p(\vec{s}\,|\,\vec{y}^{\,(n)},\Theta)\Big) \\
&= \sum_n \log(p(\vec{y}^{\,(n)}\,|\,\Theta)) + \sum_n \log\Big(\frac{\sum_{\vec{s}\in\mathcal{K}^{(n)}} p(\vec{s},\vec{y}^{\,(n)}\,|\,\Theta)}{p(\vec{y}^{\,(n)}\,|\,\Theta)}\Big) \\
&= \sum_n \log\Big(\sum_{\vec{s}\in\mathcal{K}^{(n)}} p(\vec{s},\vec{y}^{\,(n)}\,|\,\Theta)\Big)
\end{aligned}$$

which proves the claim.
□

Considering Eqn. 25 we instantly observe that $\mathcal{F}(\mathcal{K},\Theta)$ is computationally tractable if $\mathcal{K}$ is sufficiently small.

As $\mathcal{F}(\mathcal{K},\Theta)$ is a special case of $\hat{\mathcal{F}}(\mathcal{K},\hat{\Theta},\Theta)$ both free-energies are lower bounds of the likelihood $\mathcal{L}(\Theta)$. Furthermore, it can be shown that $\mathcal{F}(\mathcal{K},\hat{\Theta},\Theta)$ can be obtained as a variational lower bound of $\mathcal{F}(\mathcal{K},\Theta)$ and that the following holds.

Proposition 4
Given a generative model defined by the joint $p(\vec{s},\vec{y}\,|\,\Theta)$ let $\mathcal{L}(\Theta)$ be the likelihood in Eqn. 1 and let $\hat{\mathcal{F}}(\mathcal{K},\hat{\Theta},\Theta)$ and $\mathcal{F}(\mathcal{K},\Theta)$ be the free-energies define by Eqn. 22 and 24, respectively. Then, for all parameters $\mathcal{K}$, $\hat{\Theta}$, and Θ the following applies:

$$\mathcal{L}(\Theta) \;\geq\; \mathcal{F}(\mathcal{K},\Theta) \;\geq\; \hat{\mathcal{F}}(\mathcal{K},\hat{\Theta},\Theta)\,. \tag{29}$$

Proof
As stated before, note that $\mathcal{F}(\mathcal{K},\Theta)$ is a special case of $\hat{\mathcal{F}}(\mathcal{K},\hat{\Theta},\Theta)$ and as such a lower bound of $\mathcal{L}(\Theta)$. To show that $\mathcal{F}(\mathcal{K},\Theta) \;\geq\; \hat{\mathcal{F}}(\mathcal{K},\hat{\Theta},\Theta)$, we use Proposition 3 and apply Jensen's inequality:

$$\begin{aligned}
\mathcal{F}(\mathcal{K},\Theta) &= \sum_n \sum_{\vec{s}\in\mathcal{K}^{(n)}} \log\Big(p(\vec{s},\vec{y}^{\,(n)}\,|\,\Theta)\Big) && (30)\\
&= \sum_n \sum_{\vec{s}\in\mathcal{K}^{(n)}} \log\left(\hat{q}^{(n)}(\vec{s})\frac{p(\vec{s},\vec{y}^{\,(n)}\,|\,\Theta)}{\hat{q}^{(n)}(\vec{s})}\right) && (31)\\
&\geq \sum_n \sum_{\vec{s}\in\mathcal{K}^{(n)}} \hat{q}^{(n)}(\vec{s})\;\log\left(\frac{p(\vec{s},\vec{y}^{\,(n)}\,|\,\Theta)}{\hat{q}^{(n)}(\vec{s})}\right) && (32)
\end{aligned}$$

if $\sum_{\vec{s}\in\mathcal{K}^{(n)}} \hat{q}^{(n)}(\vec{s}) = 1$ and $\hat{q}^{(n)}(\vec{s}) \geq 0$ for all $\vec{s} \in \mathcal{K}^{(n)}$. We now define:

$$\hat{q}^{(n)}(\vec{s}) \;=\; \frac{p(\vec{s}\,|\,\vec{y}^{\,(n)},\hat{\Theta})}{\sum_{\vec{s}'\in\mathcal{K}^{(n)}} p(\vec{s}'\,|\,\vec{y}^{\,(n)},\hat{\Theta})}\,. \tag{33}$$

The choice of $\hat{q}^{(n)}(\vec{s})$ fulfills the conditions for Jensen's inequality but note that it is not a probability density on the whole state space Ω of $\vec{s}$. Inserting $\hat{q}^{(n)}(\vec{s})$ we obtain:

$$\begin{aligned}
\mathcal{F}(\mathcal{K},\Theta) \;&\geq\; \sum_n \sum_{\vec{s}\in\mathcal{K}^{(n)}} \frac{p(\vec{s}\,|\,\vec{y}^{\,(n)},\hat{\Theta})}{\sum_{\vec{s}'\in\mathcal{K}^{(n)}} p(\vec{s}'\,|\,\vec{y}^{\,(n)},\hat{\Theta})} \log\left(\frac{p(\vec{s},\vec{y}^{\,(n)}\,|\,\Theta)}{\frac{p(\vec{s}\,|\,\vec{y}^{\,(n)},\hat{\Theta})}{\sum_{\vec{s}'\in\mathcal{K}^{(n)}} p(\vec{s}'\,|\,\vec{y}^{\,(n)},\hat{\Theta})}}\right) &(34)\\
&=\; \sum_n \sum_{\vec{s}} q^{(n)}(\vec{s};\mathcal{K},\hat{\Theta}) \log\left(\frac{p(\vec{s},\vec{y}^{\,(n)}\,|\,\Theta)}{q^{(n)}(\vec{s};\mathcal{K},\hat{\Theta})}\right) &(35)\\
&=\; \sum_n \sum_{\vec{s}} q^{(n)}(\vec{s};\mathcal{K},\hat{\Theta}) \log\Big(p(\vec{s},\vec{y}^{\,(n)}\,|\,\Theta)\Big) + H(q^{(n)}(\vec{s};\mathcal{K},\hat{\Theta})) &(36)\\
&=\; \hat{\mathcal{F}}(\mathcal{K},\hat{\Theta},\Theta)\,, &(37)
\end{aligned}$$

where $q^{(n)}(\vec{s};\mathcal{K},\hat{\Theta})$ is the truncated variational distribution in (9) and where $\mathcal{F}(\mathcal{K},\hat{\Theta},\Theta)$ is the corresponding free-energy in (22).
□

Based on Proposition 4 we now know that $\mathcal{L}(\Theta)$, $\mathcal{F}(\mathcal{K},\Theta)$, $\hat{\mathcal{F}}(\mathcal{K},\hat{\Theta},\Theta)$ are a strictly ordered set.

Having established that $\hat{\mathcal{F}}(\mathcal{K},\hat{\Theta},\Theta)$ is the variational lower bound of $\mathcal{F}(\mathcal{K},\Theta)$, the following applies for the differences between $\mathcal{L}(\Theta)$, $\mathcal{F}(\mathcal{K},\Theta)$, and $\hat{\mathcal{F}}(\mathcal{K},\hat{\Theta},\Theta)$:

Corollary 1
All as above.

$$\begin{aligned}
\mathcal{L}(\Theta) - \mathcal{F}(\mathcal{K},\Theta) \;&=\; \sum_n D_{\mathrm{KL}}\big(q^{(n)}(\vec{s};\mathcal{K},\Theta), p(\vec{s}\,|\,\vec{y}^{\,(n)},\Theta)\big) \;\geq\; 0\,,\\
\mathcal{F}(\mathcal{K},\Theta) - \hat{\mathcal{F}}(\mathcal{K},\hat{\Theta},\Theta) \;&=\; \sum_n D_{\mathrm{KL}}\big(q^{(n)}(\vec{s};\mathcal{K},\hat{\Theta}), p(\vec{s}\,|\,\vec{y}^{\,(n)},\Theta)\big) - \sum_n D_{\mathrm{KL}}\big(q^{(n)}(\vec{s};\mathcal{K},\Theta), p(\vec{s}\,|\,\vec{y}^{\,(n)},\Theta)\big) \geq\; 0\,,\\
\mathcal{L}(\Theta) - \hat{\mathcal{F}}(\mathcal{K},\hat{\Theta},\Theta) \;&=\; \sum_n D_{\mathrm{KL}}\big(q^{(n)}(\vec{s};\mathcal{K},\hat{\Theta}), p(\vec{s}\,|\,\vec{y}^{\,(n)},\Theta)\big) \geq\; 0\,.
\end{aligned}$$

Proof
The first and the last equation are a direct consequence of Propositions 4 and the fact that $\mathcal{F}(\mathcal{K},\Theta)$ and $\hat{\mathcal{F}}(\mathcal{K},\hat{\Theta},\Theta)$ are both variational lower bounds of $\mathcal{L}(\Theta)$. The second equation is obtained by taking the difference of the first and the last. The fact that the difference of the KL-divergences of the second equation is greater zero, follows from Proposition 4.
□

By considering Corollary 1, we can now solve the first optimization step in (23) analytically.

Indeed, it can be shown that $\hat{\mathcal{F}}(\mathcal{K}, \hat{\Theta}, \Theta)$ is optimized w.r.t. $\hat{\Theta}$ if the values of the variational parameters $\hat{\Theta}$ are set equal to model parameters Θ:

Proposition 5
If $\hat{\mathcal{F}}(\mathcal{K}, \hat{\Theta}, \Theta)$ is the truncated free-energy of Eqn. 22, then it applies for fixed $\mathcal{K}$ and Θ that

$$\underset{\hat{\Theta}}{\operatorname{argmax}}\big\{\hat{\mathcal{F}}(\mathcal{K}, \hat{\Theta}, \Theta)\big\} \;=\; \Theta\,. \tag{38}$$

Proof
Let us first re-express $\hat{\mathcal{F}}(\mathcal{K}, \hat{\Theta}, \Theta)$ using Corollary 1:

$$\begin{aligned}
\hat{\mathcal{F}}(\mathcal{K}, \hat{\Theta}, \Theta) &= \mathcal{F}(\mathcal{K}, \Theta) + \sum_n D_{\mathrm{KL}}\big(q^{(n)}(\vec{s}; \mathcal{K}, \Theta), p(\vec{s}\,|\,\vec{y}^{\,(n)}, \Theta)\big) - \sum_n D_{\mathrm{KL}}\big(q^{(n)}(\vec{s}; \mathcal{K}, \hat{\Theta}), p(\vec{s}\,|\,\vec{y}^{\,(n)}, \Theta)\big) \\
&= \mathcal{F}(\mathcal{K}, \Theta) + \sum_n D_{\mathrm{KL}}\big(q^{(n)}(\Theta), p^{(n)}(\Theta)\big) - \sum_n D_{\mathrm{KL}}\big(q^{(n)}(\hat{\Theta}), p^{(n)}(\Theta)\big),
\end{aligned}$$

where we have abbreviated the distributions for readability. As only the last summand depends on $\hat{\Theta}$, $\hat{\mathcal{F}}(\mathcal{K}, \hat{\Theta}, \Theta)$ is maximized if $\sum_n D_{\mathrm{KL}}\big(q^{(n)}(\hat{\Theta}), p^{(n)}(\Theta)\big)$ is minimized. Now we also know from Corollary 1 that $\sum_n D_{\mathrm{KL}}\big(q^{(n)}(\hat{\Theta}), p^{(n)}(\Theta)\big) \;\geq\; \sum_n D_{\mathrm{KL}}\big(q^{(n)}(\Theta), p^{(n)}(\Theta)\big)$, where only the left-hand-side depends on $\hat{\Theta}$. Hence, the most minimal value for $\sum_n D_{\mathrm{KL}}\big(q^{(n)}(\hat{\Theta}), p^{(n)}(\Theta)\big)$ achievable is

$$\sum_n D_{\mathrm{KL}}\big(q^{(n)}(\hat{\Theta}), p^{(n)}(\Theta)\big) \;=\; \sum_n D_{\mathrm{KL}}\big(q^{(n)}(\Theta), p^{(n)}(\Theta)\big)\,. \tag{39}$$

By choosing $\hat{\Theta} = \Theta$ we can indeed satisfy the equality, and therefore know that $\sum_n D_{\mathrm{KL}}\big(q^{(n)}(\hat{\Theta}), p^{(n)}(\Theta)\big)$ takes on a global minimum there. This global minimum then implies because of (39) a global maximum of $\hat{\mathcal{F}}(\mathcal{K}, \hat{\Theta}, \Theta)$ w.r.t. $\hat{\Theta}$.
□

Note that there can potentially be other global maxima, e.g., due to permutations of the parameters without effect on $q^{(n)}(\vec{s}; \mathcal{K}, \hat{\Theta})$. Using the KL-divergences of Corollary 1 makes it salient that it is sufficient to equate $q^{(n)}(\vec{s}; \mathcal{K}, \hat{\Theta})$ and $q^{(n)}(\vec{s}; \mathcal{K}, \Theta)$, which is a weaker condition than equating $\hat{\Theta}$ and Θ. To show that $\hat{\Theta} = \Theta$ is maximizing $\mathcal{F}(\mathcal{K}, \hat{\Theta}, \Theta)$ we can, alternatively, also use Proposition 4 directly, i.e., $\mathcal{F}(\mathcal{K}, \Theta) \geq \hat{\mathcal{F}}(\mathcal{K}, \hat{\Theta}, \Theta)$. Either way, we can conclude that $\hat{\mathcal{F}}(\mathcal{K}, \hat{\Theta}, \Theta)$ is maximized for $\hat{\Theta} = \Theta$, in which case $\hat{\mathcal{F}}(\mathcal{K}, \hat{\Theta}, \Theta)$ becomes equal to $\mathcal{F}(\mathcal{K}, \Theta)$.

The results above provoke the question why $\hat{\mathcal{F}}(\mathcal{K}, \hat{\Theta}, \Theta)$ has been introduced in the first place, as $\mathcal{F}(\mathcal{K}, \Theta)$ could also be used as basis for a two stage optimization procedure, i.e., just optimizing $\mathcal{K}$ and Θ of $\mathcal{F}(\mathcal{K}, \Theta)$ in turn. While this is in principle possible, we would loose the analytical benefit that is obtained with any variational EM approximation: We would not be able to take derivatives w.r.t. Θ independently of the variational distributions, i.e., the M-step derivations (Opt 3 in Eqn. 23) would be much more intricate. By introducing $\hat{\mathcal{F}}(\mathcal{K}, \hat{\Theta}, \Theta)$, only derivatives w.r.t. the logarithm of the joint remain, reducing the analytical effort to that of standard M-step derivations. As a consequence, any M-step results of standard generative models can be used for TV-EM algorithms as it is the case for other variational approaches.

Because of the analytical solution for the first optimization step provided by Proposition 5,

the iterative TV-EM of procedure of Eqn. 23 now reduces to essentially two optimization steps:

$$\begin{array}{lrcll} \text{Opt 1} : & \hat{\Theta} & = & \Theta & \\ \text{Opt 2} : & \mathcal{K}^{\text{new}} & = & \underset{\mathcal{K}}{\text{argmax}}\big\{\mathcal{F}(\mathcal{K},\Theta)\big\} & \text{while holding } \Theta \text{ fixed} \\ \text{Opt 3} : & \Theta^{\text{new}} & = & \underset{\Theta}{\text{argmax}}\big\{\hat{\mathcal{F}}(\mathcal{K}^{\text{new}},\hat{\Theta},\Theta)\big\} & \text{while holding } \mathcal{K}^{\text{new}} \text{ and } \hat{\Theta} \text{ fixed} \end{array} \tag{40}$$

After executing Opt 1 to 3 as above, we then iterate again after setting $\mathcal{K} = \mathcal{K}^{\text{new}}$ and $\Theta = \Theta^{\text{new}}$. Because of the original optimization sequence of Eqn. 23 and the result of Proposition 5, each TV-EM step (40) increase the free-energy $\hat{\mathcal{F}}(\mathcal{K}^{\text{new}},\hat{\Theta},\Theta)$ (or leaves it constant). Notably, the same applies (after each TV-EM iteration) for the simplified free-energy $\mathcal{F}(\mathcal{K},\Theta)$ of Proposition 3. Because of Proposition 5, also Opt 2 simplifies as the compact free-energy $\mathcal{F}(\mathcal{K},\Theta)$ of Eqn. 25 can be optimized instead of the much less compact general free-energy $\hat{\mathcal{F}}(\mathcal{K}^{\text{new}},\hat{\Theta},\Theta)$ of Eqn. 22. The M-step (Opt 3) remains analog to any standard M-step of exact or variational EM such that Opt 2 is the essential new variational optimization procedure provided by TV-EM.

5 The Truncated Variational EM Algorithm

The derived theoretical results can now, following (40), be used to formulate the novel truncated variational EM (TV-EM) algorithm. Given a generative model, this meta-algorithm can directly be applied to derive a learning algorithms for the model parameters. The theoretical results of this paper guarantee that an efficient iterative procedure can be formulated that never decreases the truncated free-energy (25) which can be taken as objective function of the algorithm. None of the details of the proofs of Propositions 1 to 5 will be required, however, neither will the auxiliary free-energy (22) play a direct role. The algorithms itself is therefore very compact and any derivations required for a novel generative model will be very limited.

5.1 The TV-EM Meta-Algorithm

The TV-EM algorithm uses the truncated variational distribution (TV-distribution) $q^{(n)}(\vec{s};\mathcal{K},\Theta)$ and the truncated free-energy in the formulation of Propositions 1 to 3 and can be summarized as in Algorithm 1. The algorithm consist of an E- and an M-step. As is usual for variational approaches, the variational E-step contains a loop to optimize the variational parameters before the M-step is computed. The whole procedure is terminated if no significant changes for the parameters Θ are observed anymore.

TV-M-step. The M-step is analogous to other variational M-steps. The explicit formulation of the TV-M-step in Algorithm 1 highlights that the procedure to derive parameter update equations is directly analogous to the derivation of standard M-step equations. As only difference, the M-step equations now contain expectation values w.r.t. the truncated distributions $q^{(n)}(\vec{s};\mathcal{K},\hat{\Theta})$ in the place of expectation values w.r.t. exact posteriors, factored variational posteriors, Gaussian or

MAP approximations. Expectations w.r.t. $q^{(n)}(\vec{s}; \mathcal{K}, \hat{\Theta})$ can efficiently be computed using Eqn. 10.

Algorithm 1: Truncated Variational EM (TV-EM)

choose initial model parameters Θ and initial variational states $\mathcal{K}$

repeat

  set $\hat{\Theta} = \Theta$

  repeat — **TV-E-step**

    compute a variation $\tilde{\mathcal{K}}$ of $\mathcal{K}$

    if $\mathcal{F}(\tilde{\mathcal{K}}, \hat{\Theta}) > \mathcal{F}(\mathcal{K}, \hat{\Theta})$ **then**

      $\mathcal{K} = \tilde{\mathcal{K}}$

  until $\mathcal{F}$ *has sufficiently increased*;

  compute Θ that optimizes $\sum_{n=1}^{N} \sum_{\vec{s}} q^{(n)}(\vec{s}; \mathcal{K}, \hat{\Theta}) \log\Big(p(\vec{s}, \vec{y}^{\,(n)} \,|\, \Theta)\Big)$ — **TV-M-step**

until *parameters* Θ *have sufficiently converged*;

Where the functions $\mathcal{F}(\mathcal{K}, \Theta)$ and $q^{(n)}(\vec{s}; \mathcal{K}, \Theta)$ are given by

$$\mathcal{F}(\mathcal{K}, \Theta) = \sum_{n} \log\Big(\sum_{\vec{s} \in \mathcal{K}^{(n)}} p(\vec{s}, \vec{y}^{\,(n)} \,|\, \Theta)\Big) \text{ and } q^{(n)}(\vec{s}; \mathcal{K}, \Theta) = \frac{p(\vec{s}, \vec{y}^{\,(n)} \,|\, \Theta)}{\sum_{\vec{s}' \in \mathcal{K}^{(n)}} p(\vec{s}', \vec{y}^{\,(n)} \,|\, \Theta)} \, \delta(\vec{s} \in \mathcal{K}^{(n)}).$$

TV-E-step. The crucial difference to previous variational approaches is the truncated variational E-step. In contrast to factored variational EM, no additional functional forms with associated parameters have to be chosen. Instead, variational parameters take the form of sets of hidden states in $\mathcal{K}$. These state we term *variational states* and they are varied and selected based on the computationally tractable free-energy $\mathcal{F}(\mathcal{K}, \Theta)$. The efficiency of computing $\mathcal{F}(\mathcal{K}, \Theta)$ and the efficiency of varying the variational states $\mathcal{K}$ are the crucial procedures that determine the efficiency of the TV-EM algorithm as a whole. Any TV-EM step will, however, increase the free-energy (or will leave it constant). As it is the case for the parameters Θ of the outer loop, we can terminate the inner loop (the TV-E step) once there are no or no significant changes of $\mathcal{K}$ observed anymore. It may, however, make more sense to terminate the inner loop earlier once $\mathcal{F}(\mathcal{K}, \Theta)$ has increased sufficiently. The best strategy will depend on the concrete model and data type the TV-EM algorithm is applied to.

5.2 Variation of States

Let us finally theoretically discuss examples for varying and choosing novel states for $\mathcal{K}$ during the TV-E-step in Alg. 1.

5.2.1 Finding Variational States by Sampling

The first and easiest way to vary the states $\mathcal{K}$ in the TV-EM algorithm is to vary $\mathcal{K}$ by using new samples $\vec{s}$ drawn randomly from a suitable distribution.

Blind sampling. The most elementary way to choose variational states would be to blindly draw from a standard distribution over binary/discrete latents. Such a blind search would have the advantage of generality. However, especially for large hidden spaces the probability that new samples increase the free-energy compared to the old variational states will quickly become very small, which will make such a procedure potentially very inefficient in general. More promising would be the use of the prior distribution of a generative model to generate new variational states. With such a distribution the new states lie in the area of high probability mass if averaged over data points. The likelihood to increase the free-energy can be expected to increase faster than for fully blind search, however, large posterior mass may be located in areas of the latent space very different from areas of high prior probabilities, such that well suited variational states will again finally only be generated with relatively low probability.

Data driven sampling. Procedures that vary $\mathcal{K}$ by generating new variational states in a data driven way promise to be more efficient. For instance, if samples to vary $\mathcal{K}^{(n)}$ are drawn from the posterior distribution $p(\vec{s}\,|\,\vec{y}^{\,(n)}, \Theta)$, their joint probabilities can be expected to be relatively high (and with them the free-energy $\mathcal{F}(\mathcal{K}, \Theta)$).

5.2.2 Finding Variational States by Preselection

Truncated approximations to posterior probabilities have previously been constructed based on preselection procedures (Lücke and Eggert, 2010). Preselection meant the use of a fast estimation procedure to identify those hidden variables that could be expected to be relevant for a given data point $\vec{y}^{\,(n)}$. Only based on those variables, a set $\mathcal{K}^{(n)}$ was constructed by combining the relevant states in specific ways. Such estimation procedures made use of what was termed a *selection function* that gives high scores to latent variables relevant for a given data point $\vec{y}^{\,(n)}$. For instance for sparse coding generative models with binary latents (Henniges et al., 2010; Shelton et al., 2011, 2015b), selection functions seek to identify the basis functions $\vec{W}_h$ that have most likely generated a data point $\vec{y}^{\,(n)}$. One such example is, e.g., work by Lücke and Eggert (2010); Shelton et al. (2011) where a scalar product was used as selection function $\mathcal{S}_h(\vec{y}^{\,(n)}) \sim \vec{W}_h^{\mathrm{T}}\vec{y}^{\,(n)}$ (with $\vec{W}_h$ denoting the basis function of latent h). More generally, a selection function can be an approximation to the marginal probabilities $\mathcal{S}_h(\vec{y}^{\,(n)}) \approx p(s_h = 1\,|\,\vec{y}^{\,(n)}, \Theta)$ or may be hand-crafted for specific (possibly relatively complex) generative models (Dai and Lücke, 2014).

Motivated by these earlier works, selection functions can also be used to vary the sets $\mathcal{K}^{(n)}$ within the here derived variational TV-EM framework. Taking up the example of sparse coding with binary priors, we can, for instance, use a given selection function (e.g., the scalar product) to define a distribution for s_h to be one or zero, $q^{(n)}_{\mathrm{sel}}(s_h = 1)$, such that only hidden units deemed relevant for a data point by the selection functions are frequently activated. Such a distribution

could be defined as:

$$q^{(n)}_{\text{sel}}(s_h = 1) \;=\; \pi_h \,\frac{\mathcal{S}_h(\vec{y}^{\,(n)})}{\frac{1}{N}\sum_{n=1}^{N}\mathcal{S}_h(\vec{y}^{\,(n)})}, \tag{41}$$

which ensures that $\frac{1}{N}\sum_{n=1}^{N} q^{(n)}_{\text{sel}}(s_h = 1) = \pi_h$ as it would be required for the exact marginals $p(s_h = 1\,|\,\vec{y},\Theta)$ in sparse coding models with binary priors. The variational states resulting from such a procedure would be very similar to those that are explicitly constructed to make up sets $\mathcal{K}^{(n)}$ in earlier work. A variation of $\mathcal{K}$ based on selection functions would clearly be data driven. While it would share properties with sampling from a marginal distribution, selection functions can be more specific to the problem, e.g., by incorporating additional knowledge about the generative model and the data (e.g., sparsity). Furthermore, selection functions could be trained, e.g., by using Gaussian processes as shown by Shelton et al. (2015a) or by using any other discriminative approach (see Lücke and Eggert, 2010, for a discussion). Based on the new results presented in this paper, a criterion for a good variation could be integrated into learning procedures based on selection functions, and objective truncated free-energies could be provided.

The selection based approaches and, e.g., Eqn. 41 are specific to sparse coding generative models with binary latents. In other words, they represent a specific realization of a TV-EM algorithm. The example shows that the TV-EM algorithm is broad enough to incorporate more specific realizations but also points to more general procedures that do not necessarily rely on selection functions.

5.2.3 Other Procedures to Find Variational States

More generally, any procedure that efficiently maximizes the truncated free-energy in Alg. 1 can be used in the development for an efficient learning algorithms for a given generative model. Efficiency is hereby the crucial point, and Alg. 1 does not make this explicit. An example for an inefficient procedure consistent with Alg. 1 is, for instance, to always add new variational states to $\mathcal{K}$. According to the truncated free-energy (25) this would always ensure an increase of the free-energy but the number of states would quickly become too large. By constraining all $\mathcal{K}^{(n)}$ to a fixed and sufficiently small size, allows for maintaining efficiency. The associated general optimization procedure would than be given by:

$$\tilde{\mathcal{K}}^{(n)} \;=\; \underset{\substack{\mathcal{K}^{(n)}\\ |\mathcal{K}^{(n)}|=\text{const.}}}{\text{argmax}} \;\Big\{\log\Big(\sum_{\vec{s}\in\mathcal{K}^{(n)}} p(\vec{s},\vec{y}^{\,(n)}\,|\,\Theta)\Big)\Big\}\,. \tag{42}$$

Fulfilling Eqn. 42 implies that the criterion to increase the free-energy (25) would always be fulfilled, and any general method developed for problems with such a structure would be applicable within our framework. In general, given a data point $\vec{y}^{\,(n)}$, the goal of TV-EM is to find the set of states with the highest joint probabilities, i.e., the variational states seeks to fill all states with significant posterior mass. All mass that is not captured, makes up the difference between likelihood and free-energy as measured by the KL-divergence (19).

Even though more general than the previous approaches, Eqn. 42 is already based on a constraint of the general procedure in Alg. 1: fixed-size $\mathcal{K}^{(n)}$. Less constrained other procedures are also possible such as different sizes of $\mathcal{K}^{(n)}$ for different data points, or flexible sizes depending

on how far learning has advanced. In general, based on the results of this study, any procedure consistent with Alg. 1 that is efficient in increasing the truncated free-energy of Proposition 3 (Eqn. 25) is conceivable.

6 Discussion

In this work we have studied the theoretical foundation of truncated variational approximations to the expectation maximization (EM) approach. Our first set of results (Propositions 1 and 2) generalize the free-energy approach as introduced, e.g. by Saul et al. (1996); Neal and Hinton (1998); Jordan et al. (1999) to include discrete variational distributions with exact zeros. While this generalization is required for our study, our results are generally applicable to any discrete distribution with exact zero, i.e., they are not restricted to the specific truncated distributions of Eqn. 9 as used here. As such they generalize the standard text book derivation of variational free-energies (e.g. Murphy, 2012) to include any discrete variational distribution. Propositions 3, 4, 5 and Corollary 1 then represent results specific to the chosen form of truncated variational distributions, showing that the corresponding truncated free-energy is given by a computationally tractable expression that can be increased to (approximately) maximize the likelihood (Eqn. 25).

Relation to deterministic approximate EM procedures. Our truncated variational EM approach (TV-EM) was motivated as a novel class of variational approaches. As Gaussian variational EM is not applicable to discrete hidden variables, factored variational approaches make up the main class of comparison. Compared to fully-factored approaches (8) a main difference is that truncated approaches do not assume posterior independence. Assuming independence (no explaining-away effects) has been observed to negatively impact likelihood optimization for different types of generative models (Ilin and Valpola, 2003; MacKay, 2003; Turner and Sahani, 2011; Sheikh et al., 2014). To address such potentially harmful consequences of posterior independence, partly factored approaches (sometimes called structured variational EM approaches) have been studied (Saul and Jordan, 1996; MacKay, 2003; Bouchard-Côté and Jordan, 2009). Any deviation from fully factored approaches does, however, go along with increased analytical and computational effort, and may even provide only very little improvements compared to mean-field (e.g. Turner and Sahani, 2011). In comparison, TV-EM does not make any independence assumptions, computational tractability is instead achieved by approximating posteriors by finite sets of states. In a sense, such an approach is similar to the maximum a-posteriori (MAP) approximation which is widely used, e.g., in the domain of sparse coding (Olshausen and Field, 1996; Mairal et al., 2010) and compressive sensing (Donoho, 2006; Baraniuk, 2007). Indeed, if we constrain all state sets $\mathcal{K}^{(n)}$ to contain a single state, Alg. 1 reduces to a MAP-based approximate EM algorithm for discrete latents. As such a MAP-based procedure is a special case of TV-EM, the truncated free-energy (25) provides an objective function for MAP optimizations. Another consequence is that the variation of the sets $\mathcal{K}^{(n)}$ translates to a variation of the one-state MAP approximation. This means that we obtain with MAP a variational EM procedure that guarantees to increase a free-energy (or leaves it constant) even if an approximate MAP is changed only to increases the posterior rather than to maximize it.

Relation to a "sparse" EM algorithm. Notably, the very influential work by Neal and Hinton (1998) has never explicitly mentioned factored variational approaches. Instead the paper discussed a variational approach that shares more similarity with a truncated distribution than with a factored one. In their "sparse algorithm", Neal and Hinton (1998) suggested a variational distribution defined on a subset of the state space which only contains "plausible" values of the latents (their Sec. 5). The probabilities outside of this set were "frozen" to values of an earlier iteration but updated once in a while. The procedure there described is not efficient, e.g., for multiple-causes generative models with large state spaces because it still requires an occasional evaluation of all latent states. As Neal and Hinton (1998) discuss an application to a mixture model, this shortcoming is not very relevant for their paper. Truncated approximations do, in contrast, assume exact zero, they provide approximations without ever having to evaluate all hidden states, and describe a variational procedure for finding subspaces. However, despite these differences, the results obtained for TV-EM in this work may be regarded as connecting back to an initial (and never followed up) train of thought expressed in the work by Neal and Hinton (1998).

Relation to stochastic approximate EM procedures. While variational EM procedures are typically regarded as being deterministic, TV-EM shares many properties with stochastic, i.e., sampling based EM approaches. Like sampling approaches, we approximate probabilities (in our case posteriors) by a set of states in hidden space, and these states are then used to compute expectation values. Indeed, one option to vary the variational states in $\mathcal{K}$ for TV-EM would be to sample states from a given distribution. TV-EM also shares with sampling that the accuracy of the approximation is only limited by computational demand, in the limit of infinite computational resources both, sampling and TV-EM converge to EM with exact E-step. The same can not be said about standard variational approaches. TV-EM distinguishes itself from sampling by being a variational approach that optimizes a free-energy using variational distributions. While truncated approximations can make use of sampling as part of their optimization, also other procedures like selection functions can be applied. Also the distributions that are used to vary the states in $\mathcal{K}$ are not limited to posterior distributions while they are still potentially able to provide accurate estimates of expectation values w.r.t. posterior distributions. While these are all points of difference, the fact that both TV-EM and sampling are approximating posteriors based on finite sets of hidden states makes TV-EM the variational approximation that is most closely related to sampling. Further investigations by theoretical and empirical studies may well be motivated by such a close relations.

Relation to Expectation Truncation. Finally, the EM approximations most similar to TV-EM are given by applications of Expectation Truncation (ET; Lücke and Eggert, 2010) which has motivated this work. ET uses truncated approximations to posteriors of generative models with discrete latents. The main idea to obtain efficiently scalable approximations was hereby to preselect a small set of hidden variables and to construct small sets of states by choosing an appropriate combinatorics of this small set. While a relation of truncated approximations to variational EM was repeatedly pointed out (e.g. Lücke and Sahani, 2008; Lücke and Eggert, 2010; Sheikh et al., 2014), Expectation Truncation has never been a variational approach: the states of

$\mathcal{K}^{(n)}$ were not considered to be variational parameters, the free-energy result of Proposition 3 has not been obtained or used previously, nor the observation that a procedure like Alg. 1 is guaranteed to always increase the free-energy (or to leave it constant). As a consequence, the construction of $\mathcal{K}^{(n)}$ by preselection in place of a variational optimization according to Alg. 1 was used. Such constructions of $\mathcal{K}^{(n)}$ hereby were done using knowledge about the data and the generative model to which ET was applied to. Most frequently, learning algorithms for sparse coding like generative models were developed such that $\mathcal{K}^{(n)}$ with sparse states (many zeros) were a reasonable choice. More advanced models (e.g. Dai et al., 2013; Dai and Lücke, 2014) required more advanced and hand-tailored selection and construction procedures for $\mathcal{K}^{(n)}$. On the other hand, the previous applications to different types of generative approaches demonstrate the efficiency that can be achieved using truncated approaches (Sheikh and Lücke, 2016), showing that finite numbers of states are sufficient to obtain scalable and accurate parameter optimization algorithms. In one aspects, ET provides a result that has not been addressed here: it can be shown that optimizing the parameters of smaller generative models defined per data point approximately optimizes the parameters of the full generative model. The result, valid for a large class but not all generative models, has been used in select-and-sample approaches (Shelton et al., 2012, 2015b) and was formally proven in (Sheikh and Lücke, 2016).

Autonomous and General Purpose Inference and Learning. Other than providing a general procedure to develop a learning algorithm for a specific generative model, TV-EM may also be of relevance for the fields of autonomous machine learning and probabilistic programming, which recently got a lot of attention (Rezende and Mohamed, 2015; Tran et al., 2015; Ranganath et al., 2015). The goal of these fields of study is to provide procedures that minimize expert intervention in the generation and application of learning and inference algorithms. Typically, user intervention is required for a lot of aspects in the process of developing a concrete learning algorithm. Both standard variational approaches and standard sampling have to overcome a number of analytical and practical computational challenges. A factored variational EM approach, for instance, first has to choose a specific form for the variational distributions, and then requires a potentially significant effort to derive update equations for their variational parameters. Instead, TV-EM does not require additional analytical steps for variational E-step equations, expectation values are computed directly based on the variational states. Also in case of sampling, deriving, e.g., an efficient sampler requires additional and potentially highly non-trivial analytical work. On the other hand, the variation of $\mathcal{K}$ in TV-EM may require additional efforts. For previous truncated approximations (ET; Lücke and Eggert, 2010) the construction of $\mathcal{K}$ was very model dependent. However, given the results of this work, the previous model specific constructions could be replaced by a general purpose variation of $\mathcal{K}$. The TV-EM approach is hereby very different, e.g., from Rezende and Mohamed (2015) who use normalizing flows and consider continuous latents, from Ranganath et al. (2015) who hierarchically expands a mean-field approach to include dependencies, or work by Tran et al. (2015) based on copulas. Furthermore, Salimans et al. (2015) use Markov Chain Monte Carlo (MCMC) samplers and treat the samples as auxiliary variables within a standard variational free-energy framework. Related to results of stochastic variational inference (Hoffman et al., 2013), auxiliary distributions for Markov chains are defined and used to approximate true posteriors. Such an approach is much more indirect than TV-EM, which

directly incorporates hidden states as variational parameters. Salimans et al. (2015) also make a number of choices to define appropriate MCMC samplers and focus on models with continuous latents. The very active research on such and other very recent approaches for a "black-box" optimization framework, highlights the requirements of such approaches for advanced data models, however, and the short-comings of standard approaches. The use of collections of hidden states within variational approaches, e.g., as done by Salimans et al. (2015) or with truncated approaches (Lücke and Eggert, 2010; Shelton et al., 2011; Sheikh et al., 2014; Sheikh and Lücke, 2016), seems to be a promising general strategy in this respect.

Conclusion. Alongside of very promising recent work on improved and efficient approximate inference, more elementary approaches such as MAP approximations or basic sampling approaches remain popular choices also also for recent developments of probabilistic Machine Learning algorithms. Increasing interest in more precise approximations is motivated, however, by increasing demands for more complex data models with complex latent dependencies. Examples are novel sparse coding models or hierarchical (deep) directed models. For the purposes of this work, we have developed our results for TV-EM using a single set of latent variables $\vec{s}$. Note, however, that such a vector can contain any subsets of latents, e.g., of different hierarchical layers without changing any of our theoretical results. In practice, the model structure will have a strong influence on how concrete TV-EM algorithms work, however. For specific models and constructed variational state sets, truncated distributions have already been shown to be precise and very efficiently scalable (Henniges et al., 2014; Sheikh et al., 2014; Sheikh and Lücke, 2016). The novel theoretical results of this paper, which connect back, e.g. to (Neal and Hinton, 1998), offer mathematically grounded justification for these approaches, and, more importantly, provide a new variational framework for novel and scalable optimization procedures for advanced generative models.

Appendix: Kullback-Leibler Divergence, Likelihood, and Free-Energy

The KL-divergence is given by

$$D_{\mathrm{KL}}\big(q(\vec{s}),p(\vec{s})\big) \;=\; -\sum_{\vec{s}} q(\vec{s})\log\big(\frac{p(\vec{s})}{q(\vec{s})}\big)\,. \tag{43}$$

By taking the difference between data likelihood and free-energy we then obtain:

$$\begin{aligned}
0 \;\leq\;& L(\Theta) - F(q,\Theta)\\
=\;& \sum_n\Big(\log p(\vec{y}^{\,(n)}\,|\,\Theta) - \sum_{\vec{s}} q^{(n)}(\vec{s})\log\big(p(\vec{s},\vec{y}^{\,(n)}\,|\,\Theta)\big) + \sum_{\vec{s}} q^{(n)}(\vec{s})\log\big(q^{(n)}(\vec{s})\big)\Big)\\
=\;& \sum_n\Big(\log p(\vec{y}^{\,(n)}\,|\,\Theta) - \sum_{\vec{s}} q^{(n)}(\vec{s})\log\big(p(\vec{s}\,|\,\vec{y}^{\,(n)},\Theta)p(\vec{y}^{\,(n)}\,|\,\Theta)\big) + \sum_{\vec{s}} q^{(n)}(\vec{s})\log\big(q^{(n)}(\vec{s})\big)\Big)\\
=\;& \sum_n\Big(\log p(\vec{y}^{\,(n)}\,|\,\Theta) - \log\big(p(\vec{y}^{\,(n)}\,|\,\Theta)\big)\underbrace{\sum_{\vec{s}} q^{(n)}(\vec{s})}_{1} - \sum_{\vec{s}} q^{(n)}(\vec{s})\log\big(p(\vec{s}\,|\,\vec{y}^{\,(n)},\Theta)\big) + \sum_{\vec{s}} q^{(n)}(\vec{s})\log\big(q^{(n)}(\vec{s})\big)\Big)\\
=\;& \sum_n\Big(-\sum_{\vec{s}} q^{(n)}(\vec{s})\log\big(p(\vec{s}\,|\,\vec{y}^{\,(n)},\Theta)\big) + \sum_{\vec{s}} q^{(n)}(\vec{s})\log\big(q^{(n)}(\vec{s})\big)\Big)\\
=\;& \sum_n D_{\mathrm{KL}}\big(q^{(n)}(\vec{s}),p(\vec{s}\,|\,\vec{y}^{\,(n)},\Theta)\big)\,.
\end{aligned} \tag{44}$$

References

Attias, H. (2000). A variational bayesian framework for graphical models. *Advances in neural information processing systems*, 12(1-2):209–215.

Baraniuk, R. G. (2007). Compressive sensing. *IEEE signal processing magazine*, 24(4).

Beal, M. and Ghahramani, Z. (2003). The variational bayesian em algorithm for incomplete data: with application to scoring graphical model structures. *Bayesian statistics*, 7:453–464.

Bishop, C. (2006). *Pattern Recognition and Machine Learning*. Springer.

Bouchard-Côté, A. and Jordan, M. I. (2009). Optimization of structured mean field objectives. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, pages 67–74. AUAI Press.

Dai, Z., Exarchakis, G., and Lücke, J. (2013). What are the invariant occlusive components of image patches? a probabilistic generative approach. In *Advances in Neural Information Processing Systems*, pages 243–251.

Dai, Z. and Lücke, J. (2014). Autonomous document cleaning – a generative approach to reconstruct strongly corrupted scanned texts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(10):1950–1962.

Dempster, A. P., Laird, N. M., and Rubin, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm (with discussion). *Journal of the Royal Statistical Society B*, 39:1–38.

Donoho, D. L. (2006). Compressed sensing. *IEEE Transactions on information theory*, 52(4):1289–1306.

Ghahramani, Z. and Jordan, M. I. (1995). Learning from incomplete data. A.I. Memo No. 1509.

Haft, M., Hofman, R., and Tresp, V. (2004). Generative binary codes. *Pattern Anal Appl*, 6:269–84.

Hathaway, R. J. (1986). Another interpretation of the em algorithm for mixture distributions. *Statistics & Probability Letters*, 4(2):53–56.

Henniges, M., Puertas, G., Bornschein, J., Eggert, J., and Lücke, J. (2010). Binary sparse coding. In *Proceedings LVA/ICA*, LNCS 6365, pages 450–57. Springer.

Henniges, M., Turner, R. E., Sahani, M., Eggert, J., and Lücke, J. (2014). Efficient occlusive components analysis. *Journal of Machine Learning Research*, 15:2689–2722.

Hoffman, M. D., Blei, D. M., Wang, C., and Paisley, J. W. (2013). Stochastic variational inference. *Journal of Machine Learning Research*, 14(1):1303–1347.

Ilin, A. and Valpola, H. (2003). On the effect of the form of the posterior approximation in variational learning of ica models. In *Proceedings ICA*, pages 915–920.

Jaakkola, T. (2000). Tutorial on variational approximation methods. In Opper, M. and Saad, D., editors, *Advanced mean field methods: theory and practice*. MIT Press.

Jordan, M., Ghahramani, Z., Jaakkola, T., and Saul, L. (1999). An introduction to variational methods for graphical models. *Machine Learning*, 37:183–233.

Lee, H., Battle, A., Raina, R., and Ng, A. (2007). Efficient sparse coding algorithms. In *Advances in Neural Information Processing Systems*, volume 20, pages 801–08.

Lücke, J. and Eggert, J. (2010). Expectation truncation and the benefits of preselection in training generative models. *Journal of Machine Learning Research*, 11:2855–900.

Lücke, J. and Sahani, M. (2008). Maximal causes for non-linear component extraction. *Journal of Machine Learning Research*, 9:1227–67.

MacKay, D. J. C. (2003). *Information Theory, Inference, and Learning Algorithms*. Cambridge University Press. Available from http://www.inference.phy.cam.ac.uk/mackay/itila/.

Mairal, J., Bach, F., Ponce, J., and Sapiro, G. (2010). Online learning for matrix factorization and sparse coding. *Journal of Machine Learning Research*, 11.

Murphy, K. P. (2012). *Machine Learning: A probabilistic perspective*. MIT press.

Neal, R. and Hinton, G. (1998). A view of the EM algorithm that justifies incremental, sparse, and other variants. In Jordan, M. I., editor, *Learning in Graphical Models*. Kluwer.

Olshausen, B. and Field, D. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381:607–9.

Opper, M. and Archambeau, C. (2009). The variational gaussian approximation revisited. *Neural computation*, 21(3):786–792.

Opper, M. and Winther, O. (2005). Expectation consistent approximate inference. *Journal of Machine Learning Research*, 6(Dec):2177–2204.

Ranganath, R., Tran, D., and Blei, D. M. (2015). Hierarchical variational models. *arXiv preprint arXiv:1511.02386*.

Rezende, D. J. and Mohamed, S. (2015). Variational inference with normalizing flows. *International Conference on Machine Learning*.

Salimans, T., Kingma, D. P., Welling, M., et al. (2015). Markov chain monte carlo and variational inference: Bridging the gap. In *International Conference on Machine Learning*, pages 1218–1226.

Saul, L. K., Jaakkola, T., and Jordan, M. I. (1996). Mean field theory for sigmoid belief networks. *Journal of artificial intelligence research*, 4(1):61–76.

Saul, L. K. and Jordan, M. I. (1996). Exploiting tractable substructures in intractable networks. *Advances in neural information processing systems*, pages 486–492.

Seeger, M. (2008). Bayesian inference and optimal design for the sparse linear model. *Journal of Machine Learning Research*, 9:759–813.

Sheikh, A.-S. and Lücke, J. (2016). Select-and-sample for spike-and-slab sparse coding. In *Advances in Neural Information Processing Systems*, page accepted.

Sheikh, A.-S., Shelton, J. A., and Lücke, J. (2014). A truncated em approach for spike-and-slab sparse coding. *Journal of Machine Learning Research*, 15:2653–2687.

Shelton, J. A., Bornschein, J., Sheikh, A.-S., Berkes, P., and Lücke, J. (2011). Select and sample – a model of efficient neural inference and learning. *Advances in Neural Information Processing Systems*, 24:2618–2626.

Shelton, J. A., Gasthaus, J., Dai, Z., Lücke, J., and Gretton, A. (2015a). Gp-select: Accelerating em using adaptive subspace preselection. *arXiv*, page 1412.3411.

Shelton, J. A., Sheikh, A.-S., Bornschein, J., Sterne, P., and Lcke, J. (2015b). Nonlinear spike-and-slab sparse coding for interpretable image encoding. *PLoS ONE*, 10:e0124088.

Shelton, J. A., Sterne, P., Bornschein, J., Sheikh, A. S., and Lücke, J. (2012). Why MCA? Nonlinear sparse coding with spike-and-slab prior for neurally plausible image encoding. *Advances in Neural Information Processing Systems*, 25:2285–2293.

Tran, D., Blei, D., and Airoldi, E. M. (2015). Copula variational inference. In *Advances in Neural Information Processing Systems*, pages 3564–3572.

Turner, R. E. and Sahani, M. (2011). *Bayesian Time Series Models*, chapter Two problems with variational expectation maximisation for time-series models. Cambridge University Press.